



INCLUDES POWERFUL  
PROGRAMMING UTILITIES  
ON DISK

SECOND EDITION

# DOS

---

# UNDOCUMENTED

---

A PROGRAMMER'S GUIDE TO RESERVED  
MS-DOS® FUNCTIONS AND DATA STRUCTURES

---

*EXPANDED TO INCLUDE MS-DOS 6,  
NOVELL DOS, AND WINDOWS™ 3.1*

ANDREW SCHULMAN • RALF BROWN • DAVID MAXEY  
RAYMOND J. MICHELS • JIM KYLE

*Also available...another Documented  
Bestseller from Addison-Wesley*

# UNDOCUMENTED WINDOWS™

*A Programmer's Guide to Reserved Microsoft Windows® API Functions*

► Andrew Schulman,  
David Maxey,  
and Matt Pietrek

736 pages

Book/disk  
package with 3.5"  
1.4MB IBM disk  
0-201-60834-0

*All-inclusive coverage  
of 250 undocumented  
Windows 3.0 and 3.1 features*



*"Undocumented Windows will take a place of honor on your bookshelf."  
— PC Magazine*

Available wherever computer books are sold  
or call Addison-Wesley (1-800-358-4566)  
for more information.



# Undocumented DOS Second Edition

A Programmer's Guide to Reserved  
MS-DOS<sup>®</sup> Functions and Data Structures

---

*Andrew Schulman, Ralf Brown, David Maxey,  
Raymond J. Michels, and Jim Kyle*

---



Addison-Wesley Publishing Company

Reading, Massachusetts Menlo Park, California New York  
Don Mills, Ontario Wokingham, England Amsterdam Bonn  
Sydney Singapore Tokyo Madrid San Juan  
Paris Seoul Milan Mexico City Taipei

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book and Addison-Wesley was aware of a trademark claim, the designations have been printed in initial capital letters.

The authors and publishers have taken care in preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

### Library of Congress Cataloging-in-Publication Data

Undocumented DOS: a programmer's guide to reserved MS-DOS functions and data structures / by Andrew Schulman . . . [et. al.] — 2nd ed.

p. cm. — (Andrew Schulman programming series)

Includes index.

ISBN 0-201-63287-X

I. Operating systems (Computers) 2. MS-DOS (Computer file)

I. Schulman, Andrew. II. Series: Schulman, Andrew. Andrew Schulman programming series.

QA76.76.f63U53 1993

D05.A46 .d20

93-29037

CIP

Copyright © 1994 Andrew Schulman, Ralf Brown, David Maxey, Raymond J. Michels, and Jim Kyle

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher. Printed in the United States of America. Published simultaneously in Canada.

Chapter 1 excerpts from *Gate: How Microsoft's Agent Reinvented an Industry* by Stephen Manes and Paul Andrews. Copyright © 1993 by Stephen Manes and Paul Andrews. Used by permission of Doubleday, a division of Bantam Doubleday Dell Publishing Group, Inc.

Production Editor: Andrew Williams

Set in 10 point Galliard by Benchmark Productions, Inc.

1 2 3 4 5 6 7 8 9 BAH-9796959493

First Printing, November 1993

Addison-Wesley books are available for bulk purchases by corporations, institutions, and other organizations. For more information please contact the Corporate, Government, and Special Sales Department at (617) 944-3700 x2915.

# CONTENTS

<b>CHAPTER 1 — Undocumented DOS: The Madness Continues</b> .....	1
“Cruel Coding” and Tying Arrangements .....	3
Windows and DR DOS .....	5
Systems Rivalry and Smoking Guns .....	6
The Windows AARD Detection Code .....	7
A Countlet of Tests .....	9
A Gracious Gatekeeper .....	13
Does Beta Code Really Matter? .....	15
So What? .....	16
Microsoft’s Response .....	16
Documentation vs. Tying .....	17
Microsoft Windows Uses Undocumented DOS .....	18
WIN.COM Walks the SFT .....	20
BlockDev and INT 2Fh Function 13h .....	21
DOSMGR: Windows’ Connection to Undocumented DOS .....	21
CON CON CON CON CON .....	24
The Undocumented DOSMGR Callout API .....	24
Implementing DOSMGR Functions .....	30
Patching DOS .....	32
DOS Knows About Windows .....	34
DOSMGR and the SDA .....	35
DOSMGR and the InDOS Flag .....	35
SYSTEM.INI Settings and Undocumented DOS .....	36
KRNL386 Grows the SFT .....	37
KRNL386 and the PSP .....	37
Undocumented DOS and the Utilities Wars .....	38
Undocumented SmartDrive .....	39
Undocumented DoubleSpace .....	40
Undocumented EMN386.EXE .....	42
Microsoft Anti-Virus .....	43
No Problem? .....	45
DOS Documented .....	48
Why Leave Functionality Undocumented? .....	50
Documentation and Monopoly .....	52
Fear of Undocumented DOS .....	54
Ain’t Misbehavin’ .....	55
<b>CHAPTER 2 — Programming for Documented and Undocumented DOS: A Comparison</b> .....	59
Using Documented DOS Functions .....	60
DOS Calls from Assembly Language .....	61

**CHAPTER 2 (continued)**

DOS Calls from C .....	62
<i>int86()</i> .....	62
<i>inline Assembler</i> .....	63
Register Pseudo-Variables .....	64
DOS Library Functions .....	65
DOS Calls from Turbo Pascal .....	65
DOS Calls from BASIC .....	66
Using Undocumented DOS .....	67
Disassembling DOS .....	68
Using the Interrupt List .....	70
No Magic Numbers .....	70
Undocumented DOS Calls from Assembly Language .....	72
DOS Versions .....	74
Accessing SysVars .....	75
Undocumented DOS Calls from C .....	76
What, No Structures? .....	78
Undocumented DOS Calls from Turbo Pascal .....	82
Undocumented DOS Calls from BASIC .....	84
When Not To Use Undocumented Features .....	85
Verifying Undocumented DOS .....	86
Making Modifications .....	87
An Important Special Case: Novell NetWare .....	93
Hooking DOS: Application Wrappers .....	95
On to Protected Mode .....	99

**CHAPTER 3 — Undocumented DOS Meets Windows** 101

Calling Undocumented DOS from Windows .....	103
It Doesn't Really Work! .....	108
The Dreaded GP Fault .....	110
A DPMI Shell .....	113
Trying Out Undocumented DOS from DPMISH Programs .....	117
The Windows DOS Extenders .....	119
Inside the DOSMGR DOS Extender .....	122
How DOSMGR Handles Undocumented DOS Calls .....	126
Do Your Own XLAT .....	128
DPMI Programming .....	128
Hiding DPMI .....	131
Fixing SFTWALK .....	138
Inside the DPMI Server in VMM .....	141
Back to Windows Programming .....	143
Windows and the SFT .....	143
Walking the Device Chain .....	144
Truename .....	148
Windows and the PSP .....	151
Peeking at DOS Boxes from a Windows Program .....	155
A Brief Introduction to VxD Programming .....	173
Timing DOS Calls .....	176

<b>CHAPTER 4 — Other DOSs: From DR DOS and NetWare to MVDMS in OS/2 and Windows NT</b> .....	179
From CP/M to DR DOS to Novell DOS .....	181
The DR DOS Version Number .....	184
Undocumented Novell DOS .....	185
Watching DR DOS .....	189
Disassembling DR DOS .....	189
How Close Is DR DOS to MS-DOS? .....	190
SysVars, the Current Directory Structure, and the Redirector .....	191
The System File Tables and SHARE .....	192
Memory Control Blocks .....	192
TSRs and the Swappable Data Area .....	194
Additional DR DOS and Novell DOS Functionality .....	194
Novell NetWare .....	195
NETX and INT 21h .....	196
NetWare 4.0 and the Network Redirector .....	197
How NETX Changes INT 21h .....	198
Undocumented NetWare .....	204
OS/2 2.x: "A Better DOS Than DOS"? .....	205
MVDMS and VDOs .....	207
So What Version of DOS Is This DOS Emulation Pretending To Be? .....	208
Loading a Genuine DOS .....	210
OS/2 2.x and Undocumented DOS .....	211
New OS/2 Services for Old DOS Programs .....	213
DOS Emulation Under Windows NT .....	216
The Client/Server Model .....	217
NTVDM, NTIO, and NTDOS .....	218
Magic Pills and Bops .....	219
What Is NTVDM.EXE? .....	220
DOS 5.50 .....	222
Additional NTDOS Functionality .....	224
Undocumented NT .....	225
<b>CHAPTER 5 — INTRSPY: A Program for Exploring DOS</b> .....	229
Why a Script-Drivers, Event-Driven Debugger? .....	229
A Guided Tour .....	230
Device Drivers .....	234
Watching XMS .....	237
Dynamic Hooks .....	238
INTRSPY User's Guide .....	241
Using INTRSPY EXE .....	241
Using CMDSPY.EXE .....	241
Script Language .....	241
Syntax .....	242
INCLUDE Syntax .....	242
STRUCTURE Syntax .....	242
INTERCEPT Syntax .....	243
GENERATE Syntax .....	245
RUN Syntax .....	245
REPORT, STOP and RESTART Syntax .....	246

**CHAPTER 5 (continued)**

DEBUG Syntax	246
Predefined Constants	247
Error Messages	247
CMDSPY Compilation Messages	247
CMDSPY and INTRSPY In Operation	248
INTRSPY Utility Scripts	248
UNDOC	248
LSTOFLST	251
Log Your Machine's Activity	253
Monitoring Disk I/O	254
MEM	258
Writing a Generic Interrupt Handler	258
The Problem with Intel's INT	260
Changed Implementation in INTRSPY 2.0	261
Implementation	261
Pitfalls I Fell In	262
The Future of INTRSPY	263

**CHAPTER 6 — Disassembling DOS**

What is MS-DOS?	266
Disassembling IO.SYS and MSDOS.SYS	267
Interrupt Vectors and Chaining	271
Tracing a DOS INT 21h Call	280
Unassembling the Get/Set PSP Functions	287
Unassembling INT 21h AH=33h	288
Examining the Low-Memory Stub for DOS=HIGH	289
Examining the INT 21h Dispatch Function	291
Examining the INT 21h Dispatch Table	297
Get SysVars and the Caller's Registers	299
A Very Brief Glance at File I/O	301
Tracing a DOS INT 2Fh Call	301
How Does DEBUG Trace Through an INT?	301
INTCHAIN	302
Examining The INT 2Fh Chain	306
The MSDOS.SYS and IO.SYS INT 2Fh Handlers	308
Examining the MSDOS.SYS Handler for INT 2Fh AH=12h	309
Locating the INT 2Fh AH=12h Dispatch Table	310
Really Disassembling DOS	313
Using NICEDBG	315
Examining a Few DOS Functions	324
Examining the DOS Lseek Function	326
Other Parts of DOS	332
Am I Going to Jail for This?	333
Use the Source, Luke!	336
Microsoft's DOS OEM Adaptation Kit (OAK)	340

**CHAPTER 7 — MS-DOS Resource Management: Memory, Processes, Devices**

Memory Management	343
Memory Control Blocks	344
The HMA and UMBS	346



**CHAPTER 7 (continued)**

Making Use of UMBS	347
The High Memory Area	349
How To Find the Start of the MCB Chain	350
How To Trace the MCB Chain	351
MCB Consistency Checks	354
A More Detailed UDMEM Program	356
Allocation Precautions	364
RAM Allocation Strategies	365
<i>First-fit Strategy</i>	366
<i>Best-fit Strategy</i>	366
<i>Last-fit Strategy</i>	366
Process Management	367
Program Files and Processes	367
The COM File Format	367
The EXE File Format	367
The PSP: How It Identifies a Process	368
<i>History, Purpose, and Use</i>	369
<i>(Usually) Unique Process Identifier</i>	369
<i>Undocumented Areas of the PSP</i>	370
DOS Termination Address	371
Other PSP Fields	372
Spawning Child Processes	373
Locating Parent Processes	373
<i>Locating Ancestors</i>	373
<i>Use of this Capability</i>	373
Device Management	374
Why Device Drivers Exist	374
Hardware-Dependent Details	375
<i>Logically Required Functions</i>	375
<i>Congruence of Files and Devices</i>	375
Tracing the Driver Chain	376
<i>Organization of the Device Driver Chain</i>	377
<i>How Drivers Are Initialized</i>	377
<i>Locating the Start of the Chain</i>	378
<i>Tracing It Through</i>	378
Loading Device Drivers from the DOS Command Line	381
How DEVLOD Works	382
DEVLOD.C	385
MOVUP.ASM	394
TESTNAME.ASM	395
CO.ASM	396
Make File	396
How Well Does DEVLOD Work?	397
<b>CHAPTER 8 — The DOS File System and Network Redirector</b>	401
A Quick Overview of the System	402
The DOS File System	407
Surfaces, Tracks, and Sectors	407
Partition and Boot Records	408

**CHAPTER 8 (continued)**

The Boot Record and BIOS Parameter Block (BPB).....	410
Logical Sector Numbers and the Cluster Concept .....	412
The File Allocation Table (FAT).....	413
DOS Directory Structure.....	421
The Drive Parameter Block (DPB).....	433
Buffers and Disk Caches.....	436
The Current Directory Structure (CDS) .....	443
Contents of the CDS.....	446
Walking the CDS Array.....	448
Detecting RAM Disks.....	449
DoubleSpace Drives.....	450
Stacker Drives.....	458
Novell NetWare Drives.....	461
Manufacturing and Removing Drive Letters.....	462
System File Tables (SFTs) and Job File Table (JFT).....	465
How Many FILES?.....	469
Filename From Handle.....	472
What Files Are Now Open?.....	475
Releasing Orphaned File Handles.....	482
More File Handles.....	485
System FCBs.....	489
The SHARE Hooks.....	490
The MS-DOS Network Redirector.....	494
Using the Network Redirector Interface.....	496
<i>Front-End Hooks and Device Drivers vs. Back-End Redirectors</i> .....	497
<i>What DOS Provides</i> .....	498
<i>What a Redirector Must Supply</i> .....	503
Tracing an Open, Revisted.....	503
The Phantom.....	508
Phantom Implementation.....	509
<i>Initializing the CDS</i> .....	509
<i>The Redirector INT 2Fh Handler</i> .....	511
<i>How Do We Know the Call Is for Us?</i> .....	513
<i>Handling a Read</i> .....	515
<i>The Phantom XMS File System</i> .....	517
<i>Handling an Open</i> .....	519
<i>Handling Chdir</i> .....	522
<i>Handling Mkdir</i> .....	523
Differences Between DOS Versions.....	525
The Network Redirector Specification.....	526
Using DOS Internal Functions.....	534
The Future of the DOS File System.....	539
<b>CHAPTER 9 — Memory Resident Software: Pop-ups and Multitasking</b> .....	541
TSR: It Sounds Like a Bug, But It's a Feature.....	542
Where Does Undocumented DOS Come In?.....	544
MS-DOS TSRs.....	547
The Generic TSR.....	548

**CHAPTER 9 (continued)**

TSR Programming in Microsoft and Borland C++	550
Keeping a C Program Resident	553
Not Going Resident	555
Juggling the Stack	556
DOS Functions for TSRs	558
MS-DOS Flags	558
Get/Set PSP	560
Extended Error Information	561
Extended Break Information	563
Interrupt 28h	564
Inside the Generic TSR	565
TSR Command Line Arguments	583
Writing TSRs with the DOS Swappable Data Area	584
TSRs and Task Managers	592
Removing a TSR	601
Sample TSR Programs	604
TSR111	604
TSRMEM	605
TSRZE	607
Multitasking TSR	609
Task Switching	610
MULTI Installation	610
Timer Interrupt	610
Idle Interrupt	611
Keyboard Interrupt	611
Prints	611
Miscellaneous	611

**CHAPTER 10 — Command Interpreters**

Inside COMMAND.COM	619
Requirements for a Customized Interpreter	620
Obtaining Operator Input	622
The <i>FILE</i> Device	623
Asynchronous	623
Blocking	623
Reading asynchronously and comparing	624
Interpreting Operator Requests	626
Comparing commands with the <i>FILE</i>	626
<code>WRITE</code>	627
Command-Line Arguments and Pipes	627
Executing programs and executing programs	630
Finding and Executing Internal Commands	631
Dispatching Appropriate Processes	631
Locating and Loading External Commands	631
Dealing with <i>BAT</i> Files	632
Dealing with <i>COM</i> and <i>EXE</i> Files	632
The <i>exit</i> device	633
Installable Commands	634
<code>TSHELL</code> — a Simple Command Interpreter	638

**CHAPTER 10 (continued)**

How COMMAND.COM Works	641
Why Shells Are Their Own Parents	642
How and Why COMMAND.COM Reloads Itself	642
The Division Points	643
<i>Resident, Initial, and Transient Portions</i>	643
<i>Where These Portions Are Loaded</i>	644
Why Does +3 Sometimes Quit Working?	645
DOSKEY and How It Works	647
Using the Environment	648
<i>How COMMAND.COM Uses the Environment</i>	648
<i>Locating the Environment</i>	649
<i>Other Ways of Locating the Environment</i>	651
<i>Finding the "Active" Environment</i>	656
<i>Searching the Environment</i>	656
INT 2Eh, the Back Door	659
<i>The Function</i>	659
<i>Its Use</i>	660
Alternatives to COMMAND.COM	663
4DOS.COM	663
<i>A Total Replacement, Plus More</i>	664
<i>No Undocumented Features</i>	664
Sample Program: An Environment Editor	665
Conclusion	674
<b>APPENDIX Undocumented DOS Functions and Data Structures</b>	675
Acknowledgment	675
Sample Entry	675
<b>GLOSSARY</b>	831
<b>ANNOTATED BIBLIOGRAPHY</b>	835
<b>INDEX</b>	847

## PREFACE

This second edition of *Undocumented DOS* is radically different from the first edition, which appeared three years ago. In addition to several new chapters on the interaction between DOS and Windows, on "other DOSs," and on decompressing DOS, other chapters have been completely overhauled, not only to accommodate MS-DOS 5.0, 6.0, and Microsoft's forthcoming "Chicago" operating system, DOS 7, and Windows 4, but also to reflect what we hope is a far more sophisticated understanding of DOS internals, and of the role of a documented interface in the software industry.

### **DOS? in the 1990s??**

But what's the point of a new book on DOS? This is the 1990s, right? Windows programs are outselling DOS programs. Almost if not new development of commercial products is targeting Windows rather than plain vanilla DOS. Having produced the book *Undocumented Windows*, isn't a return to DOS, with its arcane commands and boring user interface, a form of childish regression?

Not at all. First, all those nice-looking Windows applications are really just protected-mode DOS applications. Last, 99% of the world's Windows code runs on top of MS-DOS rather than on the NT operating system. For the foreseeable future, Windows programmers will need a firm grasp of how DOS works, and of how Windows interacts with DOS.

Furthermore, plain vanilla DOS is hardly dead. DOS is still the proper tool for many tasks within the large number of DOS programs such as Microsoft's own DTPSPACE, LEX that I recently wrote from under Windows CE (we know, write on Windows, programmer, who has published already that "DOS Must Die" — nonetheless confesses to spending a major part of his work schedule at the C++ prompt. One of the major computer magazines whose editor & publisher that Windows is all important, was upgrading their PCs, producing a CD-ROM to praise Windows — using Awrite, a positively ancient DOS word processor. MS-DOS may not win any awards, and DOS-based products may have a hard time attracting attention from the novelty-crazed trade press, but the fact is that DOS remains a suitable platform for many tasks.

And yet, while the journalist we just mentioned spends most of his time at the C++ prompt, it's impossible to deny that it's a C++ prompt within a Windows/DOS box. The Windows/DOS box is a very different beast from DOS itself. One of the goals of this second edition of *Undocumented DOS* is to help software developers understand the huge changes that DOS has undergone over the last few years. Yes, DOS remains tremendously important, but in many ways it's a different DOS.

### **What's New?**

The first edition of *Undocumented DOS* appeared in October 1990, just after the release of Windows 3.0 (July 1990) and a bit before the release of MS-DOS 5.0 (June 1991). We have thoroughly updated this new edition to cover MS-DOS 5, MS-DOS 6, and Windows 3.1.

If you just want to see the changes to the DOS/Windows connection, the book *Understanding Windows* does that. Because this connection is made at a low level part of Windows, you can't see the Windows source. A primary goal of this new *Understanding DOS* is to provide readers with DOS/Windows and Windows-to-DOS interface to MS-DOS, so you can see Windows' internal structure and extension to MS-DOS, removing the illusion to be part of Windows by DOS as you capture a multitasking operating system. With DOS/Windows, MS-DOS users can see the programming and design of operating system.

*Understanding Windows* is a book for Windows users who want to see what Chapter 1 knows about the interface to DOS, including and not including, design, coding, and behavior in Windows/DOS/MS-DOS interface. For the most important user of undocumented DOS, Programmers, the connection to new Windows as a real DOS program, but makes a lot of undocumented DOS calls.

Windows 1.01 is a first step to MS-DOS's VxDs. It is a step in the way that in 1990s VxDs became the same way for PC programmers that TSRs did in 1980s. There is a book about the code mode TSRs. It is a book for programmers and TSR authors, not for the ordinary people. Some VxDs will be how programmers find the changes. A VxD book is a book more important. Chapter 1 shows the key part of DOS/Windows.

So, the book is for the readers of that is a book. The key is to find out

- The DOS/Windows connection, including DOS/MS-DOS (Chapter 1) and VxDs (Chapter 3).
- Call to understand DOS/Windows and protected-mode Windows programs (Chapter 3).
- The connection between MS-DOS and Windows because of many some problems with various Windows extensions, first the I/O and second the device abstraction, DOS and Windows are on separate Windows, so it is for a real DOS calls may contain some information on some some arrangement. It is a book to see and change, but it is not a book to see the code, it is a book to see the code (Chapter 4).
- The connection between DOS and Windows for the event of the. Rather than the connection between DOS and Windows, we try to show how DOS and Windows are really new with respect to the document by DOS/MS-DOS. Chapter 4 is a book to see the connection, DOS is a book to see the connection of the document INT 24h calls.
- The connection between DOS and Windows (see Chapter 8).
- The connection between DOS and Windows, it is a book to see the memory management and the connection between DOS and Windows (see Chapter 7).
- The connection between DOS and Windows, it is a book to see the connection between DOS and Windows, it is a book to see the connection between DOS and Windows (see Chapter 5). For example, INT 24h is a book to see the connection between DOS and Windows (see Chapter 5).
- The connection between DOS and Windows, it is a book to see the connection between DOS and Windows, it is a book to see the connection between DOS and Windows (see Chapter 8).
- The connection between DOS and Windows, it is a book to see the connection between DOS and Windows, it is a book to see the connection between DOS and Windows (see Chapter 7).
- The connection between DOS and Windows, it is a book to see the connection between DOS and Windows, it is a book to see the connection between DOS and Windows (see Chapter 9).



Excel? "Information hiding" works only when all interfaces are completely and precisely defined. But this is almost impossible! In the absence of complete and precise interface definition, you'd best bet it is to know as much about the system as you possibly can.

Secondly, a surprising and unfortunate how-often DOB programmers actually do end up needing to use some undocumented feature. Four years ago, I got involved in the first edition of this book, not because I had any particular interest in undocumented interfaces, or any awe or great figures. Microsoft's documentation department, but because in my job at Lotus I had run into situations where undocumented DOB was the only solution to a problem, and the only solution to be solved because customers were demanding it.

There may be several excellent high-end CD-ROM products, including CD Corporate, CD Investment, CD M&A, surgery and acupuncturists, and so on. These products require the Microsoft CD-ROM extension, MSCDEX. Our customers would load MSCDEX, which takes a portion of the way, solely in order to run our products, as far as this went, we carried MSCDEX as part of our products. Unfortunately, but not surprisingly, given how it works, no DOB, MSCDEX does not have a default option of automatic reboots (rebooting the machine to exit MSCDEX after running our products). They attempted to default MSCDEX to using Ken Kowler's excellent MARK and RELEASE utilities, but this did not work because MSCDEX allocates device entries, which MARK RELEASE doesn't deal with. Our customers wanted a NO-MSCDEX utility. It turned out that the only way to work NO-MSCDEX was to manipulate the DOB Current Directory Structure. As you can guess, this structure is undocumented, and so is the DOB function that lets you access it (see DRVSFC in Chapter 8).

Given how MSCDEX grades into DOB, I'm not quite sure how NO-MSCDEX, which can't control MSCDEX's changes to the CDPS, could have ever worked. In any case, it did work, and customers were happy with it. A more robust NO-MSCDEX would have involved even more use of undocumented DOB.

So, yes, it became clear that our CD Networker product, a CD-ROM network server, would have to use Microsoft's network redirector interface. Trying to get this from Microsoft proved impossible. At this point, David Maves and I realized that Lotus couldn't get this interface from Microsoft, probably because you'd rather, and it would be useful to clear an entire book on the subject of undocumented DOB. That the book turned out to be quite popular among PC programmers shows that, indeed, this material is useful. There are even a few members of programmers at Microsoft who, despite "information hiding" seem to rely on this book.

So, even this material is useful. One note of warning, however. Before incorporating an undocumented DOB call into a commercial program, always double-check and triple-check that the call is not deprecated, is a hard-to-use call, or that it exists. Even if the undocumented feature seems convenient, and if often still, it is better to use a documented feature, even if it is a little more program more code, to work in future versions of DOB. Equally important, use this price as a check on undocumented DOB: the easier it is for Microsoft to produce future versions of DOB, the easier it is to do the same thing, also makes it easier for competitors at Microsoft to surface DOB call. There are, also, a greater depth and chapter 1. Finally, use learning the undocumented way to accomplish some goal, and then figuring out a documented way, will also make you a better programmer.



## What's With All the Legal Mumbo Jumbo?

The reader expecting a straightforward technical book may be wondering why Chapters 1 and 4 contain lengthy discussions of U.S. antitrust laws, and why Chapter 6 has a whole section on the law surrounding reverse engineering. What the heck is this doing in a book for programmers?

For better or for worse, you can't view software as a narrowly defined portion of your career. In a book on undocumented programming interfaces, every now and then we have to stop staring at code and ask why a given interface isn't documented, and ask whether Microsoft might derive some "benefit" from this lack of documentation. We then have to ask the bigger question: is there a "public interest" in this arrangement? Given Microsoft's position both as the supplier of a "bottleneck" resource (the operating system) and as a key competitor for with other applications vendors that use that resource, these are not rhetorical questions. I care, wish I knew the answers to them.

These seemingly controversial issues are particularly relevant because an advance copy of some of the material in this book (early drafts of Chapters 2 and 4) were used by the U.S. Federal Trade Commission (FTC) during its three-year investigation into Microsoft. The FTC terminated its investigation with the commissioners tied 2-2 over whether to sue an injunction against Microsoft. The U.S. Department of Justice (DOJ) has now taken up the Microsoft case.

At first, I viewed the FTC investigation with deep suspicion. The issue of undocumented interfaces seemed very narrow. While I felt Microsoft had tried to "bottleneck" interfaces that were vital to other software developers, this issue didn't seem to me particularly "in the public interest." Questions from a "public interest" in the FTC case: what Microsoft might have "done" to Windows to make it incompatible with other versions of DOS, particularly Novell's DR DOS, struck me as irrelevant to the issues I discuss in Chapter 1.

But after a conference, the VARD code discussed in Chapter 1 and the Microsoft Computer SCAMP warranty code discussed in Chapter 4, I had to change my mind about a lot of this stuff. I never made up my mind about these issues as the case went on, but I was finally able to get free from the somewhat confused and contradictory discussion in Chapter 2. While I still can't decide whether Microsoft is contributing to a competitive process, or merely being super-competitive, the one thing I do know is that this is a very important question for software developers, who can't always count on their dependence on Microsoft.

After Chapter 1 went to press, and after an article on this subject appeared in *The Dr. Dobbs Journal*, I was able to meet with Brad Silverberg, Microsoft's vice president of system software, and Aaron Ravitsky, a key developer at MS-DOS, Windows 386, and Windows Enhanced mode, and a main author of the VARD code discussed in Chapter 1. Microsoft's reply appears in the December 1993 issue of *The Dr. Dobbs Journal*. Brad and Aaron did not persuade me that it was okay for Windows to be using DOS, but that Microsoft tries to document for other interfaces vendors. Nor did they persuade me of the legitimacy of the VARD code's goal of warning users of non-Microsoft versions of DOS that running Windows on that system was "untested." While Brad and Aaron were good about thinking about how to reduce tech support costs for Windows, the fact that the method they chose would unnecessarily damage Novell's reputation, with no real gain for Microsoft or its customers, seems not to have occurred to them.

I hope the reader, and the other authors, will indulge some of my attempts to take this discussion beyond the narrow confines of how to make undocumented DOS calls. The prices in the book where an assembly language disassembly listing appears directly adjacent to a discussion of some antitrust case, like *Nipote v. Chicken Delight* are my favorites. They also reflect my very mixed feelings about Microsoft's role in the software industry. Is Microsoft the steward

equivalent of Edison's Menlo Park "invention factory" or is it (as one prominent writer of programming books privately maintains) "The Dark Side of the Force"? Neither, or both.

### **How Do I Get in Touch with the Authors?**

Returning to more immediately practical issues, many readers of this book will want to know how to contact the authors with questions, corrections, suggestions, and criticisms. Many of the improvements in this second edition are the direct result of feedback from many readers of the first edition.

For reviewing comments on the book (and on Chapters 4, 6, and 8), contact  
*Andrew Schulman*  
 CompuServe 76320,302  
 Internet 76320\_302@compuserve.com

I edit a monthly "Uncommented Corner" in *Dr. Dobbs' Journal* and am always looking for new topics and new writers. I am also starting a separate *UNDOC* newsletter. For details, see the README file on the accompanying disk. If you are interesting in learning more about UNDOC's contents, price, and frequency, which I have not yet decided upon, send your name and address via normal mail, not e-mail on this one, please.

*Andrew Schulman*  
 UNDOC  
 859 Massachusetts Avenue  
 Cambridge MA 02139

This is also the address to use if you want to contact the authors, but do not have e-mail access. However, while I guarantee a response to e-mail, I can't guarantee a response to "normal" (mail) mail.

For UNDISKY (Phantom) (Chapter 5) or the network redirector section of Chapter 8, contact  
*David Muxey*  
 CompuServe 70401,305  
 Internet 70401\_305@compuserve.com

Please note that the version of Phantom on the accompanying disk has two known bugs, which I detect rarely, but not detect until the disk was already being produced. First, if DEVICE HIMEMSYS, Phantom requires BIOS HIGH (running Phantom under HIMEM without DOS HIGH hangs the system). Second, Phantom allows only two directory levels, MD1\DIR1\DIR2\DIR3 fails on a Phantom drive.

For DEVICE ENSVIDE (Chapter 7) and Chapter 10, contact  
*Jim Kite*  
 CompuServe 76703,762  
 Internet 76703\_762@compuserve.com

For the TSR library and Chapter 9, contact  
*Raymond J. Michels*  
 CompuServe 72300,2414  
 Internet 72300.2414@compuserve.com

For the Interrupt List, contact  
*Ralf Brown*  
 Internet ralf@tele.iana.pau.pa.us  
 CompuServe INTERNET: ralf@tele.iana.pau.pa.us

For updated versions of the Interrupt List, see the instructions for INTERRUPT LIST on the accompanying disk.

### **Did You Do This All By Yourself?**

Not exactly. The list of people who helped make this a better book is pretty long. The following people are not responsible for any of our errors or omissions, especially because we sometimes talk it to implement their excellent suggestions! but they are responsible for much of what is good in the second edition:

Bob Archison	Carsen Bakshold Andersen	David Andrews
Paul Andrews	Dwayne Bailey	Steve Baker
Stewart Bergman	Doug Bohig	Paul Bowers
John Brennan	Chuck Brewer	Peter Burch
Ron Burk	Barbara Cano	Frank Cavallito
David Chappell	Groff Chappell	Peter Chew
Anton Chertov	Norma D. Gilbert	F. Nicholas Cuper
John Christensen	chrtov	Ray Duncan
Grant Echols	Erck Engelke	Jon Erickson
John F. Fattah	Klaus F. F. F. F. F.	Tim Farley
Matthew Feltus	Fred F. F. F.	Klaus F. F.
Eric Fogelin	George Fulk	Isabelle Gavraud
Gils Gavraud	Steve Gibson	Drew Gidson
David R. Glanville	Barry Goldberg	John Harv
Kamer G. Hascher	Martin Heller	Robert Hammett
Tommy Ingenoso	Ian Jack	Roger Jackson
Eduardo Jacob	Dave Jewell	Phil Jollans
Steven Jones	Mike Karas	Peter Kober
Kari Kukkonen	Dmitry Kondratsev	Gene Lamb
Jim Lerner	Bill Lewis	Charlie Little
Brian Livingston	Jay Lawe	Stephen Manes
David Markin	Mike Marice	Phylis Mercer
Ted Meresla	Graham Murray	Mike O'Connor
Walt Oney	Andrew Pargeter	Kant Patel
Tim Paterson	Scott Pedigo	Jeff Peters
A. Padgett Peterson	Matt Petrek	Mike Poshnoffsky
Nathane Polish	Matthew Prea	Jeff Pruse
Dan Ranga	roeb	Art Robinson
Wendy Goldman Rubin	Tim Rowe	Sci. Rubenking
Gerald Sachs	Brett Saher	Murray Sargent

Hans Salvendy	Arne Schapers	Mitchell Schoenbrun
Peter Schütz	Larry Seltzer	Vishal Seral
Mike Shels	Dan Silber	Barry Simon
Glen Slick	Richard Smith	Kyle Sparks
Mike Spilo	John Spinks	Glenn Stephens
Al Stevens	Andrew Tannenbaum	Stuart Taylor
Rav Vardey	Frank van Golluwe	Lorraine Vaughn
Rob Walker	Jerry Watkins	Marion Westertiner
Art Wigdorn	Ben Williams	Bob Williamson
Dennis Williamson	Dan Winter	Clayton Wisloff
William J. Wonnreberger	Dan L. Wright	Manfred Young
Richard Zigler	Kelly Zyzanski	

I am sure that I must have left someone's name off this list. If so, I am sorry. There are also several employees of Microsoft whose help we would like to acknowledge by name, but can't, because the company wants to make sure that you understand that Microsoft in no way endorses this book or in any way helped with its contents. I thank you!

While everyone I've named has helped make this a better book, a few helped out to such an extent that they deserve special thanks:

Carol Chappo	John "Fritz" Frazar	Tom Earley
Gene Lands	Bill Lewis	David Markun
Larry Seltzer	John Spinks	

I hope to buy each of you several beers!

This book could not have happened without our literary agent, Claudette Moore, who is a person who makes things happen. Four years ago when I was thinking about doing a book on the low-level details of DOS, Windows, and OS/2, Claudette made the only reasonable point: that I should start out with a book called *Undocumented DOS* and that separate books of other operating environments could be written at a later time. Thanks Claudette! Thanks too to John Paul Moore, for giving this project his blessing.

At launch, my co-author, Andrew "Spks" Williams has orchestrated the transformation of this book from a messy collection of ASCII files into the attractive pages you see now. And he did a much more of this during the entire project. Spks was assisted by Jennifer Noble and in the early stages by Chris Williams. Chris arranged a special week-long undocumentation retreat to Phi Chapter Ship Inn in Scituate, Massachusetts, where large chunks of Chapters 1 and 3 were written.

Our copy editor, Meredith Roland, did a wonderful job transforming our technical code dumps into reasonable English. She is responsible for any improvement in readability which will be found by comparing what this was like in the first edition. Clearly, she did not see that last sentence.

A million thanks to Xavier Cazaot for his work on the French edition, *Les coulisses du DOS*.

At Addison-Wesley, special thanks to Phil Sotherland, Claire Horne, Lisa Roth, and Doms Ruderman. Mrs.

Thanks to my wife, Amanda Gabor, and my son, Matthew Jacob Schulman, for putting up with me the past few months.

Andrew Schulman  
Cambridge, MA

# Undocumented DOS: The Madness Continues

by Andrew Schulman

We go right at least to 120K, and some of you already own the MCB BIN too, version 6.0. MS-DOS is one of the most widely operating systems still in use. Not considerably, it is also the world's most widely used operating system. Microsoft has a lot of profitable systems, and a better DOS's lack of features, but it is not DOS's overall market accounts for its profitability and spectacular success. One estimate puts the number of commercial and enterprise desktop computer applications for MS-DOS at more than 20,000. Estimates of the installed base of DOS systems range from 80 million to 100 million. Since some of these estimates appear in marketing literature, I have been conservative and call it 75 million. You still do more such business than other operating systems. It is worth noting, but except for the few copies shipped to IBM machines, the Microsoft corporation makes perhaps an average of \$25 for each copy of the single-activated retail price copy of DOS. In 1991 alone, Microsoft's earnings from MS-DOS, according to its latest MS-DOS 5.00 "Share" statement, "What a sweet business!"

On each of these systems, or on some machines, MS-DOS provides not only its standard and core reproducible user interface of the ASCII command prompt, but also a program's interface. Just as users make DOS requests to program components via DIR \* \*.\* or COPY \* \* \* C: \ some 14-bit these days, WIN-32 programs make DOS requests to open a disk file, allocate memory, or even terminate. In doing so, a program passes into the Intel processor's 80386 register and sends an INT 21h instruction. The MS-DOS program's interface consists of several software interrupts, the most important of which is INT 21h.

Just as MS-DOS itself is undocumented, so are the documentation authors: how do programs like DOS-TOOLS, PC-TOOLS, or other DOS utilities work? Starting with a 1986 DOS program by Ray Duncan's subject, *Understanding MS-DOS*, the documentation authors have DOS programs, the standard assemblers, linker, and linker loader, the standard-sized bookstores, might say, and various different books, to help the user. No 24-bit user. At the time, only one DOS program was on the market: there really was that much to know about the operating system.

Now DOS programs, the books, and the debugger, open up to 16-bit disks and 16-bit memory allocations, and perhaps some hardware or compiled performance tweaks, contain a lengthy appendix, starting with 74 bits, to help the user. The INT 21h instruction, Interrupt Process, proceed to function 1, console input, write, read, interrupt, character output, and then, not surprisingly, to function 5, 4, 3, and 2.

Clearly, MS-DOS is a good example of a world where numerous books that are readily available are likely to be more useful than any one book. MS-DOS is very simple compared to many other computer operating systems, so it is possible to grasp DOS programming in a few days. In contrast to the unathomed depth of analysis on the system, such as UNIX, MS-DOS is an apparently small static world, in which every thing is easy to know, already is known.

Well, not quite.

Open the Microsoft official *MS DOS Programmer's Reference*, either the DOS 5.0 version from 1991 or the DOS 6.0 version from 1993, and you will find many strange gaps in the INT 21h function tables. For example, there are entries for function 51h (Get PSP Address) and function 54h (Get Vcch State), with nothing at all said about numbers 52h and 53h in between. Other quasi-official references such as *Demos: A Journey of MS DOS Programming* simply list the missing functions as "Reserved." So what are functions 52h and 53h? Two unused slots that Microsoft reserves for possible future use. Or perhaps some obscure code that Microsoft no longer uses? Maybe something useful only to DOS itself?

Not exactly. If you turn to the appendix of this book, you will find the following entries:

```
INT 21h function 52h  DOS 2+ -- Get SysVars (list of lists)
INT 21h function 53h  DOS 2+ -- Translate BIOS Parameter Block
```

Actually, you will also find functions 50h and 51h because (and recently these too were undocumented).

One of these functions (INT 21h AH 52h) returns a pointer to the SysVars internal data structure in the DOS data segment. Practically every DOS utility — existence, whether written by Microsoft or by a third party (DOS editors such as Central Point, Norton, Quantum, Quarterdeck, Star Electronics, and so on) — uses function 52h and SysVars (also known as the List of Lists). Yet the function really is undocumented. Microsoft uses it a little, major third party editors vendors use it, yet Microsoft has never acknowledged its existence. This is not a "test program" or the monopolization of a vital resource. We try to answer this question in the course of this chapter.

There is just one more crucial hole in the programmer's interface to MS DOS. There is a similar gap between functions 54h and 56h. This gap hides a crucial bit of functionality: function 55h (Create PSP). Why? DOS or Windows start a program, they use this function. Many third party extensions to DOS (such as third party x86 DOS Extenders and NT LEMs) or Microsoft's Windows for Workgroups (WfW) Microsoft removed NT LEM from MS DOS 6.0 at the last minute — use function 55h too.

Another hidden area of DOS is function 50h. The *MS DOS Programmer's Reference* describes function 50h as follows: the inevitable question of whether there are other functions such as 5000h, 5100h, and so on. Not surprisingly, there are. In contrast to what the Microsoft official reference indicates, function 50h consists of 12 sub-functions that handle an assortment of tasks, including DOS call back network, Net, Function Call, and support for DOS recalcitrancy (Get Swappable Data Area). In fact, the one with function 50h Microsoft does document — in certain aspects that don't make much sense, unless you know about the other ones.

Besides INT 21h (DOS internal DOS software) interrupts such as INT 21h which contain entire undocumented subsystems, the most important of these is the Network Redirector (INT 21h function 10h). This function is vital to even to many non-programmers who use DOS networking software, yet Microsoft has never documented this function's programming interface, despite years of repeated requests from developers.

There are other numerous, in fact, the most apparent portion of undocumented DOS. The real core of undocumented DOS is its data structures: the DOS internal variable table (SysVars), the System Call Table (SCT), the Current Directory Structure (CDS), and numerous other structures that this book describes in detail.

Although MS DOS is a sea of code, it is nonetheless far from being a self-enclosed static world like a fractal equation whose graph is infinitely complex. This small piece of code called DOS contains many unexplained areas. Unexplained, at least, by Microsoft.

## “Cruel Coding” and Tying Arrangements

As a concrete example of how a documented DOS is used in the real world, consider Microsoft’s enormously successful Windows operating environment, which runs on top of MS-DOS. Later in this chapter, we will see that Windows (especially Enhanced mode) totally depends on our legal knowledge of undocumented DOS functions and data structures, in particular INT 21h function 52h and SysVars. That Windows relies on DOS in an undocumented DOS is a use of interesting interest to explore in a book on undocumented DOS. In doing this, we may have a wide reputation.

The computer trade press has given a lot of attention to the question of whether Microsoft’s position as an the leading supplier of PC operating systems has given its applications an unfair advantage over applications from other vendors. More of the readers searching for a good book *Undocumented Windows* had to do with the fact that Microsoft applications were created assuming undocumented Windows functions. But the computer trade press has also been focused on the wrong way. It is equally important to ask whether Microsoft’s software operating system supplier gives an unfair advantage by bringing out *new* and *open* operating systems, such as Windows and Windows NT.

The point will be made throughout this chapter that even though we call the “part of the operating system” and hence Microsoft’s natural domain to do with is this section, in fact, DOS and DOS extensions are some of the most fiercely competitive areas in the PC software marketplace. The existence of a large third party market for new utilities and extensions shows that a large group of consumers considers the DOS separate from the operating system. The question of competing in DOS market whether Microsoft is abusing its monopoly position as nearly sole supplier of the world’s pre-dominant operating system, and whether it is also taking it leaves from this position, is a different one.

We can inquire further the question by looking at the relationship between Windows and MS-DOS. The DOS Windows connection raises important issues for Microsoft. One issue is, should a company must avoid what is called a “tying arrangement” — such as using one product, such as Windows, requires that someone another product, such as MS-DOS, from the same company, and prevents you in some way from buying a competitor’s compatible product, such as Novell’s DR-DOS. Tying arrangements come under U.S. antitrust law, particularly section 7 of the Clayton Act, which explicitly prohibits a firm competitor the use of sales or commodity “in a contract, agreement, or understanding that is less than purchase, contract, sale, lease or other deal of the goods” — not a competitor or competitors of the seller or seller, the effect “may be to substantially lessen competition or tend to create a monopoly.”

On the other hand, Microsoft very much wants to show that to promote its Windows system it to upgrade *whenever* is the latest release of MS-DOS. For example, the October 29, 1992 Microsoft “Reviewer’s Guide” for DOS 6 states:

### *MS-DOS must be a specific platform for Windows*

Windows has become a standard. More than 100,000 PCs are shipped with Windows, and the market is growing. The Windows application market is approaching the size of the MS-DOS applications market. Given the widespread use of Windows, we plan to evolve MS-DOS over time to:

- Provide the base technology that Windows needs to improve.
- Become more fully integrated with Windows.

Likewise, the RT-ADM WRITING included with Windows 3.1 states:

### *Running Windows within Operating System Other than MS-DOS*

Microsoft Windows and MS-DOS work together as an integrated system. They were designed together and extensively tested together on a wide variety of computers and hardware configurations. Running Windows version 3.1 on an operating system other than MS-DOS can, of course, unexpected results or poor performance.





## Windows and DR DOS

During the abortive three year U.S. Federal Trade Commission (FTC) investigation into Microsoft (now taken up by the U.S. Justice Department) attorneys from the FTC were carefully looking at the relationship between MS DOS and Windows. For example, *PC Week* (November 30, 1992), reporting on the FTC investigation noted that

Another area of contention is whether Microsoft modified Windows to prevent Novell's DR DOS operating system from running with the graphical environment. DR DOS requires updates to work with both Windows 3.0 and 3.1.

Certainly there is a history of incompatibilities between Windows and DR DOS. But it is vital to filter out which of these incompatibilities are *not* Microsoft's fault. There is a vague, widespread notion that every incompatibility between Windows and DR DOS somehow reflects an attempt by Microsoft to crush its competitors. The following letter to the editor of *PC Magazine* (September 15, 1992) is a good expression of this idea:

In response to Microsoft's cruel coding, Windows 3.1 and DR DOS 6.0 are incompatible. A little message about the MS DOS extender popped up whenever I typed "win" at the prompt in DR DOS. Luckily, update disks came from Digital Research and solved all my problems.

While it is unclear what "cruel coding" means exactly, the general idea is evidently that Microsoft introduces code into MS DOS and/or Windows for the sole purpose of blocking its competitors, especially Novell. As we will see, this is not such a crazy idea. But it does not follow that every incompatibility between Windows and DR DOS therefore reflects back on Microsoft.

In this particular case, Microsoft is not at fault. The actual message that appears is "Standard Mode Load in MS-DOS Extender. This message comes from the Windows component DOSX.FIL when you are trying to run Windows 3.1. Standard mode on DR-DOS 6.0 it does not occur in Windows 3.1 End user mode. The term "E.P." here indicates a violation of one of the rules of protected mode, it usually indicates a bug in a program. Indeed, I find from one of the DR DOS developers reveals that in fact, this one was DR DOS's problem.

On the subject of Windows interaction with DOS, DR DOS 6.0 did not function correctly with Windows 3.1, however, this was not a fault to be attributed entirely to code made by Windows. The problem was caused by a compatibility fix in a TNC transfer one came to the "NE" test of task base, the setting of the TNC call as required by some old NEC processors. Windows DOS X extender introduced a few system errors with switching to protected mode and hang!

Windows 3.0 and 3.1 do make a number of error-prone updates to DOS if the version of DOS is 5.0 or higher. We are still prior to Novell's DOS 7.0, pretending to be version 6.0MP. At 3.31 so Windows does not see these updates under DOS 3.31. Windows is not as dirty as most networks, so it does not cause serious problems.

So much of case in this case, for cruel coding. Yet the FTC's Bureau of Competition took the position of Microsoft "cruel coding" very seriously. For one of its attorneys asked, is specifically a sort what Microsoft might have "put" in Windows to "take it deliberately incompatible with DR DOS.

At first, this idea sounds completely ridiculous, and is reminiscent of the old, oft-repeated but never proved story that MS DOS 3.0 contained code to deliberately break Lotus 1-2-3. According to mythology, Microsoft's slogan was "DOS will do, till Lotus don't do it!" see, for example, Charles Bergerson and Charles Morris, *Computer Wars*, p. 66. No one has ever shown that DOS 2.0 in fact contained such code or even explained what it would look like. Perhaps

```
if (exe_file == "123.EXE") then
  halt_and_catch_fire();
```



The two crucial words here are "sole" and "artificial." Surely, Microsoft should be allowed to improve Windows even if we exist at night, literally, in the DR DOS. This is a genuine part of the competitive process, and virtually anyone (judiciary included) has rights, even access to the courts. For example, many manufacturers of so-called "plug-compatibles" tried—in 1970s to have the courts characterize IBM's System 360 as "predatory innovation." The courts rejected these claims along with eventually also IBM's IBM 9600 S2, thereby effectively creating precedents to prove that the defendants' design had no reflecting value to consumers. (Stephen J. Ross, *Principles of Antitrust Law*, 1993. Technically, "consumer's" is a word in italics.)

A fascinating article on this subject is "Predator: Systems Rivalry," A Report by James Crawford, Alan Wake, and Robert Wing, *Consumer Law Review*, June 1983, who describe "systems rivalry" as follows:

Suppose that company A manufactures a product system with two components, A1 and A2, each sold separately. Company A has monopolized power over A1, but company B competes in the market for the second component, B, its competitors being C, D, E. They consumers initially choose a product system composed of either A1 and A2 or A1 and B2. Company A now introduces a new product system, A1 and A2, which serves roughly the same function for consumers as the old product system. Component B2, however, is incompatible with A1. Furthermore, company A discontinues the sale of A1 or else replaces A1 substantially (injection molding). As a consequence, consumers switch to the new product system, and company B is driven from the market for component two. When, if ever, should the antitrust laws sanction company A for driving B out of the market?

Here's a great comparison: here to Windows A1, MS DOS, A2, and DR DOS, B2. This scenario, which has nothing to do with operating system software, addresses one aspect that there's really little else in the question set to which Microsoft.

Normally, A driving B out of the market would be competition, not a abuse, that's the goal of competition and should be protected. So, if you tell what this case to be "anti-competition" and becomes predatory. "Predatory Systems Rivalry" provides a good test:

the plaintiff must be a "bona fide" producer in the same relevant market; the magnitude of R&D expenses by a preponderance of the evidence; the plaintiff would have incurred a "smoking gun" or document or an admission that clearly reveals the plaintiff's culpable state of mind; and, most of the R&D decisions. Alternatives: the plaintiff would prevail if the innovation in question yields such benefits to consumers that no case could be made that it had anticompetitive, but an anticompetitive purpose.

This is what the question of the DR DOS Windows competition comes down to. Has Microsoft done anything to Windows and/or MS DOS for an anticompetitive purpose. Let's look at the closest thing we have to a "smoking gun."

### **The Windows AARD Detection Code**

If you were one of the 15,000 Windows 3.1 beta testers, and if you happened to be using DR DOS rather than MS DOS, you probably encountered the following, seemingly innocuous, yet odd, error message generated by WIN.COM:

```
Non-fatal error detected, error #2726
Please contact windows 3.1 beta support
Press ENTER to exit or C to continue
```

As we'll see, this message is a visible manifestation of a piece of code in Windows whose implementation is technical, a support artifact, and which extensively uses undocumented DOS APIs.



Yes, WIN.COM contains this code, which looks like something out of a teenage virus writer's nightmare. The code in Figure 1-1 would disable any attempt to set breakpoints in DEBUG. In Figure 1-1, we were unable to use DEBUG to examine this DEBUG-defeating code, because rather than set breakpoints, we used the DEBUG assembler's command. Unfortunately, WIN.COM's redefinition of INT 3 does not affect modern debuggers such as SoftMega's SoftICE, which runs the debugger and debugger in separate address spaces. In any case, these attempts to throw purveyors off the track are in themselves quite revealing.

These attempts succeed at more effective later on, where the code actually *disassembles* WIN.COM, passes the disassembled code through a NOP filter before execution, thereby, for a large amount of what we can only call "reasoning." You can see this for yourself by using DEBUG to unassemble the Windows 3.11 main program of WIN.COM on the SYSTEM\WIN31\\*.exe (dated March 10, 1992) with Windows 3.11 SE1. It goes to 0040:WIN.COM, and DEBUG's disassembled code starts at offset 3C12h and ends at 3EAFh.

```
C:\WIN31>type aard.scr
u 3ce2 3fae
n
C:\WIN31>debug win.com < aard.scr > aard.lst
```

Earning the resulting AARD.LST by itself is something reasonable when a product causes you many hours of grief.

For some reason, while most of the whole AARD assembly file contains plain text, Microsoft copy got stuck and the bytes AARD and RSVV probably the program's signature. This is a fueled speculation on the guess, see, for example, *Windows 3.11: The Inside Story*. While the FCC contains its Senses About Microsoft's Messes at *Inside*, August 1995, that contains an article, I think, other than Aaron Rembold, one of Microsoft's most highly skilled and well-respected programmers and a member of the staff of both MS-DOS and Windows 3.11, called this block of code has become known as "the AARD code."

**A Gauntlet of Tests** Code 1, as written, and shown on a small link to refer to key out the AARD code (Figure 1-2) shows a pseudocode summary of the assembler code. I can see clearly AARD code, which is really a part of Windows 3.11, and should separate a lot. MS-DOS uses unimplemented APIs here, and structures to check for genuine MS-DOS or PC-DOS.

### Figure 1-2: Pseudocode of the AARD Code in WIN.COM

```
move [and fixup] code from 2D19h to 4EDh
call code at 4EDh
  call AARD code at 39B2h
  -- see below
  IF (AX doesn't match 2000h)
    AND IF (control byte is non-zero) .. added in retail
    THEN overwrite BYTE at 4EDh to a RET instruction
***
IF (byte at 4EDh is a RET instruction)
  THEN issue non-fatal error message
call AARD code at 39B2h
  print INT 1, 2, 3 and at invalid code to confuse debuggers
  call undocumented INT 21h AH 52h to get SysVars (list of lists)
  copy first 30h bytes of SysVars to stack
  copy first 4 bytes (OPB ptr) of copy of SysVars to stack
  IF DOS version >= 10.0 (i.e., OS/2)
    THEN don't set [bp+190h], so eventually OR AX, 2000h fails
ELSE
  check fields in SysVars to ensure non-zero:
```

```

SysVars[0] -- Disk Parameter Block (DPB)
SysVars[4] -- System File Table (SFT)
SysVars[8] -- Clock device
SysVars[12h] -- Buffers header
SysVars[16h] -- Current Directory Structure (CDS)
SysVars[20h] -- CON device
SysVars[22h] -- Device driver chain (MUL device next ptr)
If no SysVars fields are zero (MS-DOS, or WIN.COM in DR DOS)
  THEN set [bp+196h] so that eventually OR AX, 2000h succeeds
Else some are zero e.g., HIMER STS in DR DOS)
  THEN don't set [bp+196h], so eventually OR AX, 2000h fails
copy code
jump to copied code
copy and XOR code
jump to copied and XORed code
; the following crucial part was figured out by Geoff Chappell
If a redirector is running (INT 2Fh AX=1100h)
  AND if default upper case map (int 23h AH=38h) in DOS
    data segment (undocumented INT 2Fh AX=1203h)
OR If no redirector
  AND if first System FCB header (SysVars[1Ah]) offset == 0
  THEN DOS is considered okay
Else (e.g., WIN.COM, SMARTDR.EXE, etc. in DR DOS)
  int h clear part of [bp+196h] so eventually OR AX, 2000h fails
restore previous INT 1, 2, 3
jump back to saved return address

```

As seen in Figure 1-2, the AARD code relies heavily on undocumented DOS functions and data structures, several of which are undocumented. The function 57h is, for example, the LFN entry in SysVars, which is populated with the LFNs of the current directory. SysVars contains pointers to other DOS internal data structures, such as the current file structure (CDS) and system file table (SFT). The AARD code is a combination of these pointers in SysVars, assuming that they are valid.

As a reader will expect, DOS normally does not implement these internal DOS data structures. For example, the MS-DOS files and folders are implemented through the initial set of tests. However, the AARD code is designed to pass further. We look at the way we would construct a good what we will see after we have finished our research on the undocumented DOS data structures.

It is important to note that code was originally made to check such as HIME.MSYS if files under DR DOS. In fact, the code was based on DR DOS. It is possible to obtain a genuine CDS and don't set up a genuine CDS, but the code will not work. Thus, the Windows 3.1 beta HIME.MSYS will not work. The code will only work if DR DOS is installed. Similarly, the AARD code runs under Windows NT, but it will check the CDS pointer. Finally, the code identifies OS 2.03 as Windows 3.11, but it will check the CDS pointer. So, for example, OS 2.2.1 masquerades as OS 2.03. This is not an appropriate, the most disguised test, however, appears at the end of the AARD code, which is highlighted in Figure 1-3.

### Figure 1-3 The Crucial AARD Test for DOS Legitimacy

```

If a redirector is running (INT 2Fh AX=1100h)
  AND if default upper case map (int 23h AH=38h) in DOS
    data segment (undocumented INT 2Fh AX=1203h)
OR If no red factor
  AND if first System FCB header (SysVars[1Ah]) offset == 0
  THEN DOS is considered okay

```

As we will see, however, a check for the DOS code is also in the book *DOS Internals*. It checks to see whether a chosen redirector, such as MMIOX, is running. In this case, AARD checks that DOS is not a system file, such as a CDS data segment, which is located by calling and executing DOS INT 2Fh AX=1204h, see below, appears. If no redirector is running, AARD checks that the pointer to the first System File Control Block (FCB) is located on a paragraph boundary, that is, it

has a 0 offset. AARD gets the pointer to the first System FCB from offset FCB at SysVars which is retrieved earlier using undocumented DOS function 52h. All versions of MS-DOS pass this test — a version of DR-DOS does.

To test whether this interpretation of the encrypted and heavily obfuscated code is correct, Listing 1-1 shows MSDETECT.C. This program performs the same tests as the original AARD code, but without the obfuscations and with more informative "error" messages. MSDETECT.C succeeds under all versions of MS-DOS (MS-DOS 3.31, MS-DOS 5.0, MS-DOS 6.0) and fails under all versions of DR-DOS tested (DR-DOS 5.0, DR-DOS 6.0, StarView, DR-DOS 7.0). It is scanning under DR-DOS with a redirector. MSDETECT.C fails with the message "Critical case: SysVars[0] FCB data segment!" Otherwise, it fails under DR-DOS with the message "First System FCB not located from paragraph boundary!"

### Listing 1-1: MSDETECT.C

```
/*
MSDETECT.C
Andrew Schulman, May 1993
From Undocumented DOS, 2nd edition (Addison-Wesley, 1993)
See also Dr. Dobb's Journal, September 1993

A duplication of Microsoft's MS-DOS detection code from Windows 3.1
WIN.COM (WIN.CHX), SMARTDRV.EXE, HIMEM.SYS, SETUP.EXE, MSD.EXE

The original Microsoft code (with the initials 'AARD') is heavily XOR
encrypted and obfuscated. Here the encryptions and obfuscations have
been removed. The AARD code also attempts to disable debuggers by
pointing INT 1, 2, and 3 at invalid code (FFh FFh), this also has
been avoided here.

The original AARD code non-fatal error message has been replaced
with more informative messages about the exact problem (i.e.,
difference from MS-DOS).

Geoff Chappell deciphered the original code's tests (upper case map
segment, System FCB) in the case where the preliminary SysVars tests
fail. Some of this material is discussed in Geoff's forthcoming book,
"BOS Internals" (Addison-Wesley, 1993).

Optional command line switches DOREDIR and NOREDIR have been added
to aid testing both crucial tests whether or not a redirector is
actually installed.

Build program with Microsoft C: cl msdetect.c
*/
#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <dos.h>

typedef int BOOL;
typedef unsigned char BYTE;
typedef unsigned short WORD;
typedef unsigned long DWORD;
typedef void far *FP;

BYTE far *_dos_getsysvars(void),
FP _dos_getcasemap(void),
WORD _dos_getdataseg(void),
BOOL _dos_isredirector(void),

void fail(const char *s) { puts(s); exit(1); }

main(int argc, char *argv[])
{
    BYTE far *sysvars,
    int do_redir = 1,
```

```

// this is just to make testing easier
// can fake presence or absence of a network redirector
// to exercise both tests
if (argc > 1 && argv[1][0] == '-')
{
    char *s = strdup(argv[1][1]);
    if (strcmp(s, "DOREDIR") == 0)
        do_redir = 1;
    else if (strcmp(s, "NOREDIR") == 0)
        do_redir = 0;
    else
        fail("usage: msdetect [-doredir] [-noredir]");
}

if ((sysvars = __dos_getsysvars()) == 0)
    fail("INT 21h AE5200h returns 0");

if (_osmajor >= 0xD0)
    fail("DOS version >= 10, this is OS/2 (or early NT beta)");

#define SYSVARS ofs          (*(FP far *) &sysvars[ofs])
#define SYSVARS_TEST(ofs, msg) if (! SYSVARS(ofs)) fail(msg)

SYSVARS_TEST(0, "Disk Parameter Block (DPB) pointer is 0");
SYSVARS_TEST(4, "System File Table (SFT) pointer is 0");
SYSVARS_TEST(8, "CLOCKS device pointer is 0");
SYSVARS_TEST(0x12, "BUFFERS header pointer is 0");
// next test fails in device driver init (HIREM SYS) under DR DOS 6.0
SYSVARS_TEST(0x16, "Current Directory Struct (CDS) ptr is 0");
SYSVARS_TEST(0x0C, "CON device pointer is 0");
SYSVARS_TEST(0x22, "Device chain ptr (from AUL) is 0");

if (! dos_isredirector()) && (do_redir == -1)
    do_redir = 0;

// the following tests fail under DR DOS 5 and 6
// (and under beta of Novell DOS 7)
if (do_redir)
{
    FP casemap,
    puts("Doing redirector test");
    casemap = __dos_getcasemap();
    if (FP_SEG(casemap) != __dos_getdataseg())
        fail("Default case map isn't in DOS data segment");
    printf("case map @ %iip\n", casemap);
}
else
{
    puts("Doing no-redirector test");
    if (FP_OFF(SYSVARS(0x1A)) != 0)
        fail("First System FCB not located on paragraph boundary");
    printf("System FCB ptr @ %iip -> %iip\n, sysvars=0x16, SYSVARS(0x1A));
}

// if get here, everything checks out
puts("All tests check out - must be MS-DOS");
return 0;
}

// undocumented function
BYTE far *__dos_getsysvars(void)
{
    // could initialize ES:BX to 0 0 but the MS code doesn't do this
    _asm mov ax, 5200h
    _asm int 21h
    _asm mov dx, es
}

```





points to `Recv` that in real mode, a single memory location can be addressed by different pointers, there are many combinations of different segment and offset values that all resolve to the same physical address. It is therefore equivalent. Not surprisingly, Windows is unaffected by this change to these pointers. As you will find in [Figure 1-4](#), the only software that noticed was `MSDOS EXEC` and the `AARD` code in `WIN.COM`.

**Figure 1-4: Making the AARD Code Fail**

```
C:\>DEBUG CHAP1>syndeb
Microsoft Symbolic Debug Utility
Windows Version 3.00
(C) Copyright Microsoft Corp 1984-1990
Processor is I80386

... The first System FCB is stored in this configuration at 0116:0040,
... so demormalize the pointer at that location, changing it from
... 05E4:0000 to 05E4:0040. This points to the same exact location,
... but since the offset isn't zero the AARD test fails.
-dd 0116:0040 0040
0116:0040 05E4.0000
ad 0116:0040 05E0 0040

... Now demormalize the pointer for the default case map. I had to
... disassemble the code for INT 21h AH 3Bh to find where this is
... located. The pointer is stored here at 0116:12AB. Below, the
... pointer is changed from 0116:0C15 to 01E5:0005. This points
... to the same exact location, but the segment isn't 0116 (DOS data
... segment) anymore, so the AARD test fails.
-dd 0116:12AB 12AB
0116:12AB 0116 0C15
-ud 0116:12AB 01E5 0005
-0

C:\>WIN30>win
Non-fatal error detected: error #2726
Please contact Windows 3.1 beta support
Press ENTER to exit or C to continue
C:\>DEBUG CHAP1>syndeb
Default case map isn't in DOS data segment!
```

These results are mostly surprising. Even *Recall My Last Known Drive* article "Replace Case Mapping Function with File Access" ([Q49736](#), April 29, 1993) reminds us that it is perfectly legal and documented for a DOS program to hook `INT 21h AH 3Bh` and replace its built-in case mapping function.

It has recently been noted about half a dozen times, Windows relies heavily on undocumented `IOPL` calls. This is not a problem, which we will take up later, but perhaps the `AARD` code is simply verifying that the internal `IOPL` data structure supports the undocumented `IOPL` calls upon which Windows relies. If this is the case, part of the `AARD` code does test for the presence of internal `IOPL` data structure. This would be a reasonable component, since Windows is commercial. *If the AARD code stopped there, it would be a good idea to compare your test code.* However, DR `IOPL` does not fail these tests in `WIN.COM`. DR `IOPL` has some weaknesses. System FCB case map test in [Figure 1-3](#). And Windows does not rely on that behavior at all.

In other contexts, such as `MSDOS`, if you modify the operating system, it would be perfectly acceptable to expect the `IOPL` data structures to see that they were the same as would be expected with a stock `MSDOS`. One of the `AARD` related searches for those `IOPL` internal data structures on which Windows actually relies that too could be a genuine, though we still have to ask why Windows wouldn't expect that `MSDOS` that Microsoft chooses to disclose to the rest of the software developer community. However, the fact that `WIN.COM` and other programs incorporating `AARD` code



**So What?** A non-fatal error message in a beta version—that's it. If you have an axe to grind with Microsoft, you may have expected something more nakedly robber-baronesque behavior. If this is the worst that can be found, perhaps things are not so bad after all. None, however, that at least one other "strike against" has appeared: the amazing, warrants-related error message in QuickC and Microsoft C 6.0, discussed in Chapter 8. While it is difficult to second-guess the precise goal of the encrypted and obfuscated AARD code, its results are clear enough: Windows beta sites that used DR DOS rather than MS DOS would have been scared, if not using DR DOS. "Doctor, every time I do this I get a non-fatal warning." "Then stop doing it!"

Speculating further, a conceivable goal of the AARD code was to delay Novell's attempts to get out a commercial DR DOS that is comparable with Windows 3.1. Recall that the AARD code was added to Windows from the start of the beta test—the scenario runs as follows: Novell notices the message and queries that its customers will get weird errors whenever they run Windows 3.1. Its logical next step is to try to reverse-engineer WIN.COM in order to figure out why. However, the code has been cleverly written to obscure the precise test that is being conducted—and you thought I was "Obfuscated C" contests served no practical purpose!

Effect of the AARD code is to create a slow and highly artificial test of DOS compatibility. Encryptions and obfuscations make it difficult for a competitor even to determine what is being tested. A indication that the AARD code's obfuscation was successful is the fact that a beta version of Novell's latest version of DR DOS fails the test, even though this version is otherwise far more compatible with MS DOS than previous versions.

Novell is hard to believe that Microsoft would exert such efforts to clutter something with as little market share and economic problems as DR DOS. DR DOS's best place in the sun was done by the fact that MS DOS 4.0 and PC DOS 4.0 displaced IBM, and as soon as MS DOS 5.0 came out any reasonable use DR DOS quickly disappeared for retail customers at least. But consider OEM pur-chasers, who may purchase up to 1000 copies of DOS on every machine they sell. For every single non-IBM PC word sold, Microsoft extracts a substantial licensing fee. See, further, see Magee and Andrews, *Gate*, pp. 264-265, for a detailed discussion of Microsoft's per-machine (rather than per-copy OEM pricing) of MS DOS. Microsoft must hold Intel's experience with AMD and Cyrix, fears any threat one matter how lame, to this wonderful monopoly position.

**Microsoft's Response** While Microsoft has provided no formal response to these accusations, Brad Silverberg, Microsoft's Visual Systems Software, was quoted in *Updata* (August 1993) as providing the following advice for the AARD error message: "Windows is designed for MS DOS. If DR DOS is 100 percent compatible with MS DOS, as this case implies, then it will never show up."

The implication is that if DR DOS gets an error message from Windows that MS DOS doesn't get, then the limitation is Novell's fault, and proof that DR DOS isn't 100% compatible—which it is. It remains unclear, since see Chapter 8, the problem with this argument is that, as seen in Figure 1.4, the AARD code, even for DOS compatibility, is 100 percent artificial. By Microsoft's definition, any MS DOS user that gets a warning for being identical with MS DOS—and therefore, a violation of copyright—is "100 percent DOS compatible."

As for why Microsoft checked Lusk's output in the final shipping version of Windows 3.1, Silverberg told reporter Wendy Roth: "It wasn't worth the hassle." Roth: "The hassle of?" Silverberg: "Oh, people like you asking me questions like this." *Updata* (August 1993).

Finally, a high-level manager at Microsoft was reportedly told us that there can't be any malicious intentions by the AARD code, because Microsoft is "agnostic" regarding DR DOS. This seems unlikely given the effort required to write this tricky code. Its presence in five otherwise unrelated programs also points—fairly convincingly, that, as it is unlikely that five so different programs are all maintained by the same person. In fact, the programs probably fall under the domain of several different product managers or divisions. Finally, this same Microsoft official has argued that the AARD code represented

a legitimate effort by Microsoft to ensure compatibility and reduce technical support costs for Windows. The AARD code, in other words, is really a form of quality control to protect Windows' "good will." The argument goes something like this:

- Windows is not just one DOS application, but part of a "suite" or "platform" system, i.e., MS-DOS. It is six ways to be, just another way of saying that Windows relies heavily on undocumented DOS internals.
- Microsoft has no guarantee that other vendors' versions of DOS support all the DOS functionality that Windows requires.
- Microsoft is not obligated to test Windows on these other vendors' versions of DOS.
- Therefore, Microsoft needed a way for Windows to test for the presence of some other, untested, version of DOS and to warn the user of a *potential* problem, i.e., that the combination of Windows with this version of DOS was untested and was not guaranteed to work.
- The test had to be encrypted and obfuscated in some way, or else, if the user DOS vendor could easily circumvent the test by changing its DOS combination, wouldn't the code, was testing for, without doing anything else, make its DOS "not compatible"?

Note that this defense of the AARD code rests on Windows' use of unencrypted DOS calls. According to the argument, the AARD is not deliberately incompatible. That is, it serves a legitimate purpose for consumers, because Windows depends so heavily on undocumented DOS calls and structures that Microsoft has no guarantee that Windows will work on any other DOS. In other words, Windows tests on undocumented DOS calls, and the AARD code is a *necessary* attempt to shore up this shaky foundation.

But in this case, why not test for the specific undocumented DOS functionality on which Windows relies? Why test for features that Windows doesn't care about, like constructing the file extension redirector, System I/O, or interrupt? Microsoft has constructed Windows to avoid any deliberate incompatibility. The Windows developers tried to create a working environment on DR-DOS only because DR-DOS failed to pass their own strict internal quality test. This does not show any formal quality control or the user protection of Windows' "good will."

**Documentation vs. Tying** If Microsoft needs to exert control over how Windows runs on other platforms, perhaps systems search their minds for the "right" way to do it, i.e., search by tying Windows to MS-DOS, or via other mechanisms and copyrighted code. If Windows runs on some undocumented aspects of DOS, why Microsoft can't *document* what's under the hood, why does it not underlie DOS. Instead of the AARD code, Microsoft could have issued a specification defining the internal DOS calls and file structures that Windows requires. If DR-DOS fails to meet these specifications, so what? Or Microsoft could produce a PSL or other code that MS-DOS may provide a slow, inefficient, equivalent to the undocumented calls similar to IOCTL\_IOCTL in Windows. The alternative to Microsoft's tying of Windows and MS-DOS, or other code, is *documentation*.

Interestingly, when a company says that two of its products must be tied in some way, the courts have long held that the company can't instead produce specific rules of minimum quality standards. For example, in *Seagraves, Candler, Adams, 1971*, where Chickadee Design Company refused to purchase all their supplies and fixtures directly from Chickadee Design, court held that the company's "exclusive" is "affecting a certain control over the address of its specifications, the size of its, concrete machinery, and the lip and space moves." Similarly, Microsoft could document the "secret" calls and specs, as it were, Windows expects of the underlying DOS. Without this competition for DOS, a not a "level playing field."

If, however, every one agrees that companies should be required to disclose tying arrangements with documentation. For example, Robert Berk, a well-known judge writing many cases, pointed

house of it. A Chicago school of economists—in a section of his *The Inefficient Paradigm* (1978) devoted to “Technological Interdependence or the ‘Protection of Goodwill’”—says that “the manufacturer is likely to understand the technical problems of his machines better than his users” and that “the writing of specifications and the contract problems required to make sure they are complied with is also very likely to be more expensive than supplying the related good oneself.”

But even if by the “transaction costs” are too high for proper documentation and specifications proper to faces are too expensive, and the production of technically interdependent products is best handled internally from “supplying the related good oneself” as a normal matter rather than leaving it to the market. So much for the free market, which is apparently too inefficient! It follows from this, of course, that so-called “Chinese Walls” between the different divisions of such a firm are a very easy idea to come up with, but that use of an undocumented language is a more subtle one.

For there need to be “Chinese Walls” not only because Microsoft keeps claiming they exist, when in fact they do not. Microsoft uses its monopoly position as operating systems supplier to leverage its competitive position as applications vendor, not just because this is good *commercial practice*. These nonverbal “Chinese Walls” are also allowing more than just software engineering’s “firewalls” (narrow and well-documented interfaces). Despite the claims that non-arrangements are a very uncompetitive alternative to proper documentation and public interfaces, engineering experience tells us otherwise.

Is it right for Microsoft to be able to use undoc-oriented DOS calls in its Windows product? Is Microsoft obliged to either document all of these calls or to erect a “Chinese Wall” (i.e. a well-defined public interface) between DOS and Windows. On the one hand, these are both pieces of common software. Windows is not a few thousand lines of code, it perhaps a Solidate engine, an application. On the other hand, Windows competes on the desktop (as well as in the GUI) with desktop PCs from other companies. Microsoft certainly has the work sites on top of DOS by both incorporating knowledge of DOS internally into Windows and by making DOS more Windows-compatible. No other company has that, but why should it? Should Microsoft have this luxury or should it be required to spell out exactly how its products fit together?

## Microsoft Windows Uses Undocumented DOS

Let’s examine the DOS/Windows connection in greater detail. To properly cover this subject would require a fairly large separate book, we know because we have already written several hundred pages already and there are many more to be written. So, certainly, we can’t get some idea of what DOS and Windows are, or see a list of the various Windows calls to a DOS program that you start by typing WIN at the DOS prompt or just simply by pointing WIN at your AUTOEXEC.BAT file. And a Windows system will certainly run if programs and drivers for WIN.COM kick off.

But suppose that protected-mode multitasking, graphical windows, dynamic linking, device independence, Windows operating system, and a whole host of other features like single tasking, real mode MS-DOS operation, even more tricks to know than requires, and surprisingly intimate knowledge of how MS-DOS works. It takes knowledge that Microsoft uses to implement Windows but if you just do not want to be a vendor.

Let’s look at exactly how undoc-oriented DOS calls Windows makes. How do you find out what DOS functions are needed to get a program going? If you had access to the source code, you could just look at a chapter in this series, *Assembly*, but there is an easier way if you’re not interested in what DOS calls a program makes. The architecture of MS-DOS lets you look into system interrupts, including INT 21H-based. Who can write a routine that hooks INT 21H and other DOS interrupts, and that tells you whenever a program makes an undocumented DOS call.

David Morse designed the program INTR.Y for this very purpose. It is an event driven, single-byte DOS keyboard, but you can also use for many tasks that have nothing to do with undocumented DOS. Chapter 5 of this book describes a new enhanced version of this program in detail. You

can write an INTRSPY script that logs information to a file every time a program makes any undocumented DOS call. A very simple INTRSPY script that monitors some undocumented DOS calls (but that doesn't use many INTRSPY features) appears in Listing 1-2. This script also catches only a few undocumented DOS calls. For example, this script logs memory creation (INT31) and file system functions because Windows generates so many of these functions that they "overflow" from the INTRSPY results buffer. Windows makes many other undocumented DOS calls besides those that UNDOC.SCR traps, but the few that UNDOC.SCR does capture still prove to be fruitful.

It is important to understand here what we're discussing here is undocumented DOS calls made by Windows. This is sometimes a separate subject from that discussed in *Undocumented Windows*. That book covered the system calls that Windows (and early Windows) use to handle the control-level subject of which Microsoft applications use these calls. That book did not, however, point out Windows itself uses on top of DOS or when undocumented DOS functionality Windows relies upon.

### Listing 1-2: INTRSPY Script UNDOC.SCR

```

;;
;; UNDOC.SCR
;;
Intercept 21h
function 52h on_exit output 2152 Get List of Lists ES BX
function 55h on_entry output 2155 Create PSP DX SI
function 53h on_exit output 2153 Translate BPB
function 50h subfunction 06h
on_exit output 215006 Get DOS$SWAP DS SI
function 60h
on_entry output 2160; Canon File: " (DS:SI->byte,ascii,64)
on_exit same line -> (ES:DI->byte,ascii,64)
;;
;; Use the next functions and ints 20h and 27h to show which
;; program made the undoc DOS call, and to show termination
;;
function 46h
on_entry
output "-----"
output (DS:DX->byte,ascii,64)
function 54h on_entry output "-----"
function 31h on_entry output "-- TSR -----"
Intercept 20h on_entry output "--"
Intercept 27h on_entry output "----- TSR -----"
;;
;; Too many int 2Ah to show. Could use INTRSPY counters though.
;;
Intercept 2fh
function 13h
on_entry output "2F13 Set Disk Handler " DS " " DX
on_exit same line - (Prev: " DS " " DX " )

```

Figure 1-6 shows output from INTRSPY when starting Windows 3.1 Enhanced mode under MS-DOS 5.0.

### Figure 1-6: Selected Windows 3.1 Enhanced Mode Undocumented DOS calls

```

C:\WIN31\WIN.COM
2152: Get List of Lists: 01% 0026
2152: Get List of Lists: 01% 0026
-----
C:\WIN31\system\win386.exe
2F13: Set Disk Handler 338A:1AC5 (Prev F000:9C13)
2F13: Set Disk Handler F000:9C13 (Prev 338A:1AC5)
2F13: Set Disk Handler 338A:0ABB (Prev F000:9C13)
2F13: Set Disk Handler F000:9C13 (Prev 338A:0ABB)

```

```

2152: Get List of Lists: 0116:0026
2152: Get List of Lists: 0116:0026
2F13: Set Disk Handler: 3384:00B7 (Prev: F000:9C13)
2F15: Set Disk Handler: F000:9C13 (Prev: 3384:00B7)
      tons of 2F13 calls
2152: Get List of Lists: 0116:0026
215006: Get DOSMAP: 0116:0320

```

```

C:\WIN31\system\IRRN386.EXE
2152: Get List of Lists: 0116:0026
2155: Create PSP: 5576, 0100

```

There, too, we see that when starting Windows in Enhanced mode, WIN.COM runs WIN386.VI, WIN386.EXE, a collection of Windows virtual device drivers (VxDs). A few key VxDs are the Virtual Memory Manager (VMM), the MS DOS Manager (DOSMGR), the MS DOS Network Manager (DOSNET), the Virtual 8086 Mode Memory Manager (V86MMGR), the Virtual mouse and Keyboard Controller Device (VPIFD), and the Virtual Block Device (BlockVxD).

For the purposes of describing the interaction between DOS and Windows, the key VxD is, as you might expect, the VMM (VMMGR). When trying to understand why Windows calls the functions shown in Figure 14, DOSMGR is where we'd spend most of our time. It isn't long enough to say that Windows calls lots of unadorned DOS calls because there really isn't any such single call in a Windows-based Windows's production of any separate programs, libraries, drivers, and initialization files; it is important to point out adorned DOS calls made by Windows down to specific users.

Another VxD, SHLH, handles KRNL386.EXE, which is a Windows DLL containing the kernel portion of the Windows application programming interface (API). For example, KRNL386 contains the Windows desktop's kernel, which knows how to load Windows executables, Windows device drivers, and the network driver (NetBIOS). KRNL386 handles other DLLs like USER and USER32, the INTRSPV suspension support calls to the DOS low-level function INT 21h AH=33h, the kernel's use of device drivers being accessed, setting Windows file (FILE DRV) and Network File System (NFS) programs running in the SYSTEM32 shell—setting things up. By default, the Windows Program Manager, just as COMMAND.COM sets the only possible window (MS-DOS PROGRAM.MAN.EXE) as the only possible interface for the Windows desktop.

In Figure 14, too, for the Windows, on earlier versions parts of Windows calls INT 21h API=70h SetVars, INT 21h API=80h Create PSP, INT 21h AX=5000h Get SDA, and INT 21h API=50h Set Disk Handler. As noted earlier, Windows makes many additional undocumented DOS calls, but the few shown here are enough.

Let's briefly try to see why Windows makes these calls. INTRSPV only shows *what* to see *why*, you need either to disassemble a success of Windows with a product such as Windows Source from A Communications Company or Windows under a debugger such as Soft ICE from Nu-Mega.

### WIN.COM Walks the SFT

Why does WIN.COM? In a well-documented Get SysVars function (INT 21h AH=52h). This function is responsible for setting the internal SysVars structure within the DOS data segment. This structure is the contrast point to many other key internal DOS data structures, such as the Memory Control Block (MCB), the System File Table (SFT) chain, the Current Directory Structure (CDS) chain, and so on. For this reason, this structure is sometimes called the List of Lists.

WIN.COM kernel contains one different call, INT 21h AH=52h. Some of these are only necessary for the system's existence, such as when you run Windows in standard mode; others, in our opinion, call INT 21h AH=52h we saw earlier in the VxD code, occur every time someone types WIN. As one example, in disassembly with WIN.COM shows that WIN.COM uses the SysVars table to load the List SFT. WIN.COM then walks the SFT chain, counting the number of available files, to



ensure that you have at least FILES=30 on your system. Figure 1-7 shows disassembly of this code. If this code looks strange, you may want to read Chapter 7 and then come back here.

**Figure 1-7: WIN.COM Walks the SFT Chain**

```

4C40 182B      mov ah,52h
4C40 182A      int 21h                ; get sysvars into ES:BX
4C40 182C      xor ax,ax              ; init number of files = 0
4C40 182E      les bx,dword ptr es [bx+4] ; sysvars[4] ptr to first SFT
4C40 1832      DO_NEXT_SFT:
4C40 1832      add ax,es:[bx+4]      ; SFT[4] is number of files
4C40 1836      cmp word ptr es [bx],0FFFFh ; SFT[0] ptr to next, at end?
4C40 183A      je short DONE         ; if so, done
4C40 183C      les bx,dword ptr es [bx] ; if not, walk linked list
4C40 183F      jmp short DO_NEXT_SFT
4C40 1841      DONE:
4C40 1841      sub ax,16h           ; done walking list of SFTs
4C40 1844      jnc short loc_ret_010 ; 16h = 30
                          ; Complain if not FILES=30+

```

This code is very similar to the SETBACK example in Chapter 8. In one sense, it is one of undocumented DOS's amazing deal-it-almost-seems-trivial. You might ask, is this what all this "undocumented" is all about. The answer is yes. Many of the ways Microsoft designed undocumented DOS seem trivial in the same way.

But think for a moment: your favorite word gets the FILES=30 as a *documented* way, while only a few nonusers know of the CONFIDENTIAL way and read up the code. But that method is totally unreliable. A user could have booted the system with a low FILES number, then changed C:\CONFIG.SYS, thought about booting up, and instead of looking down Windows Manual, used the undocumented DOS and Windows—let's say, I keep it—code cover-litre paper change that can only make a big difference from the stability of a program. Note: This is not to imply that Windows is not stable.

### BlockDev and INT 2Fh Function 13h

Back in Figure 1-6, the first set of calls from WIN386 is to INT 2Fh AH=E8h, which, as noted in the appendix, sets the DOS Set Disk Interrupt Handler from the BlockDev VxD. A true WIN386 is making this call. BlockDev is new to Windows 3.1, replacing the Windows 3.0's first hard disk device. The FastDisk support in 3.1 simply provides 32-bit disk access, even on BlockDev. BlockDev needs to hook the 445 disk interrupt, INT 13h, but it can't say that the interrupt doesn't cut disk access and make that hard disk look like INT 2Fh or other error. Instead, it hooks INT 13h hook *in its own way*. Via INT 2Fh AH=E8h, DOS provides a way to reactively hook INT 13h's interrupt function. BlockDev can get its hooks in before any other Win386's handlers.

For a more complete discussion of BlockDev, FastDisk, and INT 2Fh function 5Bh, see Chapter 1 of Geoff Chappell's book *DOS Internals*.

### DOSMGR: Windows' Connection to Undocumented DOS

The other WIN386 undocumented DOS calls—Figure 1-6 are coming from the DOSMGR.VXD. This is where we get to the heart of Windows' connection to undocumented DOS. Just a few more calls appear—Figure 1-6. For some, these functions return pointers to structures in the DOS data segment; the functions open up a vast amount of DOS internals to Windows.

In the *Initialize* routine (Figure 1-8), you can see that DOSMGR stores the return address in EAX, 21h AH=57h in several different ways. After calling INT 2Fh AH=52h, DOSMGR makes *init* and *via an Exec* function provided by XMM, see Chapter 7. The EAX register provides a handle, as to get the value of DOS's data segment. DOSMGR stores this value in a variable we will call *DOS\_DS*. DOSMGR also shifts the value left by four to form a 32-bit hex address: *DOS\_DS \* 16*. To this, it adds BX to form a 32-bit linear address to the SysVars table, *SYSVARS + 16*.





If you've used tools such as the QEMM programs EFILES.COM and FASTDRV.COM, resize and move DOS memory structures like the SEI and CDS. The whole point is to run EFILES or FASTDRV with LOADLHE. Chapter 2 presents FASTDRV, a clone of Quarterdeck's FASTDRV. These programs work by updating the pointers in SysVars. But since DOSMGR only checks SysVars once, DOSMGR can't set up for the SEI or CDS pointers to change, and if it isn't prepared for the game, LOADLHE within SysVars to change either. So, while you can run EFILES or FASTDRV before Windows loads, running them from within an Enhanced mode DOS box is a recipe for disaster. This section works technical reviewers writes: "Only an idiot would expect that to work." These pointers in SysVars are change-only Enhanced mode is running, which means that all VMs have separate copies—hence SysVars pointers. We'll take advantage of this fact in Chapter 3, in a Windows program, ENUMDRV, that displays the CDS in every VM; see Listing 3.27.

DOSMGR's complete dependence on the information it gets from SysVars, using pointers to the CDS, SEI, device count, and so on, in so many places that it's difficult to summarize. The use of multiple copies of CDSs or DOSMGR is a bit so trivial. The single call that DOSMGR makes to INT 19h allows it to specify a large amount of DOS internals to DOSMGR, making it possible for DOSMGR to look at multiple DOS boxes relatively happily on top of a single copy of DOS.

### CON CON CON CON

If we're looking for the tool DOSMGR needs to manipulate DOS internals is available in SysVars. For example, while SysVars contains a pointer to the SEI chain, and DOSMGR can determine the number of SEI entries by asking the SEI to compute as WIN.COM, LoadLHE in figure 1.7, SysVars does not provide the size of an SEI entry. DOSMGR needs to know the size of an SEI entry. Unfortunately, this size changes from one DOS version to the next; see Chapter 8 and the Appendix. Rather than using a lookup table to give the size of an SEI entry for each supported DOS version, an SEI entry is 386 bytes in DOS 3.0, 3.1, and 3.11, 500 bytes in 3.1, 3.3, and 3.86 bytes in 3.0. DOSMGR figures out the size of an SEI entry by CON five times, searches the first 512k of memory for the string CON, and then repeats the search to find the next occurrence of CON, and the next, and so on. DOSMGR then repeats the search for CON to find the next size of an SEI entry.

KRN=386 and KRN=286 use this same found technique as a part of *grazing* the SEI to get more information. This technique is from, because it is so easy for something to go wrong. Geoff Chappert reports that Error was a fun task by creating a device driver with the string "CON CON CON CON CON" and DOSMGR and KRN=386 found the string, and thereafter assumed that SEI entries were four bytes wide!

Memory managers such as QEMM or EMM386 may contain entries to allow the SEI entries to upper memory. If you're possible to overcome this, because DOSMGR won't find the string of CONs in the first 512k of memory, these entries can keep Windows from loading. If before starting Windows you've set a string that moves the SEI, Windows 3.1 can fail with the message "Unsupported DOS version" as though it were using DOS before that time, and it will not use SEIs.

To work around the problem of memory managers relocating the SEI tables to upper memory, *PC Magazine*'s "Somebody's SEI Entries" (January 17, 1993) published a batch file that creates a string of CONs 386 bytes apart in lower-level memory. You might wonder why DOSMGR doesn't end up manipulating these simulated SEI entries instead of the genuine ones. Fortunately, DOSMGR uses the five CONs just to get the size of an SEI entry, and then uses function 52b to find the entries themselves.

### The Undocumented DOSMGR Callout API

A more important example of how DOSMGR must sometimes supplement its use of undocumented DOS structures is to return to the CDS table. To use the table, DOSMGR must know the size of an individual CDS entry for a single drive. Recall that to instance the CDS, DOSMGR must

pass CDS 11h to `ADD` is attached together with the size in bytes of the CDS entry. `INT 21h AX=1607` (MS-DOS 5.0) knows the number of CDS entries, but it doesn't know how many bytes each entry is.

In most instances that require a table of DOS data structures, you'd like to know how big they are. For example, the `MS-DOS` program in Chapter 2 uses the DOS version number to determine whether a CDS entry is 516 bytes or 536 bytes. I'd like perhaps that this had some way to do this, but and that it leaves Windows owners a more intimate knowledge of DOS internals. `DOSMGR` uses another mechanism to find the size of a CDS entry: its supplement to use of undocumented DOS. `DOSMGR` has a private interface to DOS 5.0 and higher.

As each `VxD` loads, it causes a `VxD` initialization broadcast. `INT 21h AX=1607` with its `VxD` ID number in `DX`. `VxDs` that take broadcasts or that provide services to other programs use device ID numbers as well. For Microsoft `VMM` (for example), its `VxD ID` is `1` and `DOSMGR` has `VxD number 15h`. What Microsoft's Windows Device Driver Kit (WDK) `3.11 and DOS: Adaptive Guide` `VxDs` clearly documents the `INT 21h AX=1607` `VxD` initialization broadcast. Microsoft says absolutely nothing about the actual broadcast or API data, and supports 5.0 or 6.0 `VxDs` built into Windows. In other words, Microsoft doesn't even concern generic `INT 21h` (it claims not to want to find out what notifications or APIs a piece of Windows software `DOSMGR` provides, the `DDK` is silent).

Microsoft sets the `DOSMGR VxD` callout `INT 21h AX=1607` `DX=0` as a private method between MS-DOS and Microsoft Windows. A Microsoft internal document (APIs, `Kernel`, `MS-DOS Instance Data`) indicates how to use this interface, but it is not publicly available; the company apparently it was briefly available in beta versions of the Quick Help files for the Windows 3.1 `DDK`. Actually, it is an API not just for getting instance data from DOS, but for the `INT 21h` `DX=0` even without Microsoft's explicit documentation; several programmers have discovered `DOSMGR` and MS-DOS to see what they have to say to each other.

This `DOSMGR VxD` callout `INT 21h AX=1607` `DX=15h` provides several subfunctions specified in the `CX` register. Table 1-1 lists these subfunctions. For further details, see the appendix entry for `INT 21h AX=1607` as well as Geoff Chappell's book `DOS Internals`.

**Table 1-1. DOSMGR VxD Callout Subfunctions**

- 0 Query instance processing, get patch table
- 1 Set patches in DOS
- 2 Remove patches in DOS
- 3 Get size of DOS internal data structures
- 4 Query instance data structures
- 5 Get device driver size

`DOSMGR` uses these calls (see, for example, `cdh` `0` `0x00000011` and `0x00000012`) and `MINDOS` systems `INT 21h` and responds to these calls. Yes, MS-DOS 5.0 and 6.0 call back to you about Windows. For some of the functions, `DOS` expects a signature track `CDX` `AX=0` (check "signature" of `A2M1097` for the signature of word shadows). DOS 5.0 and 6.0 only implement subfunctions 0, 3, and 4. The implementation of the function 4 in DOS 5.0 and 6.0 serves to guarantee that no DOS data was inconsistent in this interface. On the other hand, should the implementation of MS-DOS-like behavior seem to be correct, and indicate that I have taken care of installing DOS internal data structures itself, `DOSMGR` is ready for its mission in Figure 1-11.

**Figure 1-11: DOSMGR Determines if MS-DOS Does Instancing**

```

05B3C mov [ebp,Client_AX],1607h ; VxD callout
05B42 mov [ebp,Client_BX],15h ; 15h = DOSMGR
05B48 mov [ebp,Client_CX],4 ; subfunc 4 = query inst data struct
05B4E mov eax,2fh
05B53 VMICal Exec Int ; do INT 2fh

```





```

#pragma pack(1)
typedef struct {
  #ifdef __TURBOC__
    unsigned short bp,di,si,ds,es,dx,cx,bx,ax;
  #else
    unsigned short es,ds,di,si,bp,sp,bx,dx,cx,ax; /* same as PUSHA */
  #endif
  unsigned short ip,cx,flags,
} REG_PARAMS,

volatile unsigned dosmgr_calls = 0;
int pass_through[10] = { 0 }; /* list of DOSMGR calls ok to pass to DOS */
#define DOSMGR_MAGIC_DX 0xA2AB
#define DOSMGR_MAGIC_AX 0xB97C
void interrupt far int2f(REG_PARAMS r);
void finterrupt far *old)(void);
void fail(char *s) { puts(s); exit(1); }
main(int argc, char *argv[])
{
  int i, func;
  if (argc < 2) fail("usage: nodosmgr [win] <func list to pass through>");

  for (i=2; i<argc; i++)
    if ((func = atoi(argv[i])) < 10)
    {
      pass_through[func]++,
      argv[i] = (void *) 0,
    }

  old = (void (*)(interrupt far *))(void) dos_getvect(0x2F); // hook 2F
  dos_setvect(0x2F, int2f);
  spawnvp(P_WAIT, argv[1], &argv[1]); // run command (e.g., Windows)
  _dos_setvect(0x2F, old); // unhook 2F
  printf(" %u calls to DOSMGR 2F/1607/15 interface\n", dosmgr_calls);
  return 0;
}

void interrupt far int2f(REG_PARAMS r)
{
  if ((r.ax == 0x1607) && (r.bx == 0x15))
  {
    dosmgr_calls++;
    if (pass_through[r.cx])
    {
      if (r.cx == 1)
      {
        r.bx = r.dx; // do just what DOS does
        r.dx = DOSMGR_MAGIC_DX; // but don't pass down to DOS
        r.ax = DOSMGR_MAGIC_AX;
      }
      else
        _chain_intr(old);
    }
    // otherwise, don't pass down to DOS
  }
  else
    _chain_intr(old);
}

```

When `NOJDOSMGR.WIN` is run on MS-DOS 6.0, Windows starts to come up, but then Windows aborts with the message "ERROR: Unsupported MS-DOS version." This is precisely the message that DOS 5.0 or higher workalikes get if they fail to implement the DOSMGR interface.



NODOSMGR has an option to pass through a specified DOSMGR subfunction. For example, running NODOSMGR WIN 0 bypasses all subfunctions except C/A 40. Query instance processing (see Table 1-1) which is passed down to DOS. Likewise, through subfunctions 0, 2, 3, 4, or 5 and results in an "ERROR: Unsupported MS-DOS version" failure from Windows DOSMGR.

However, running NODOSMGR WIN 1 has an interesting effect. Windows may "filter" other words at least when returning on generic MS-DOS 5.00 or 6.00 only DOSMGR call subfunction 1. Set patches in DOS see Table 1-1 seems to be absolutely necessary. And MS-DOS's implementation of this function is trivial so avoid in fact that NODOSMGR does it even better passing this function down to DOS. As you can see in Listing 1-3, NODOSMGR 1, each sets BX to 134 and puts the NODOSMGR magic signature in DX:AX. This is all MS-DOS does in its hand-off to INT 2Fh AX=1607h BX=15h C/A=1.

```

loc_4678
FOCB:4678  8B0A          MOV  BX,DX
FOCB:467A  EB0F          JMP  loc_4688
;
;
loc_4688:
FOCB:4688  8B7CB9       MOV  AX,B97C
FOCB:468E  8A4BA2       MOV  DX,A24B

```

So, simply by passing back a few magic numbers, we can convert the "ERROR: Unsupported MS-DOS version" message and allow Windows to run. At least conceptually, the resulting Windows DOS configuration is unstable, we're using it now to write this chapter, and fix a few DOS boxes open along with Click and Software.

Of course, this is on top of a copy of genuine MS-DOS. In fact, subfunction 1 is not sufficient. If subfunction 0 is not passed down to MS-DOS, DOSMGR fails back to its own methods for patching DOS, see Chapter 1, *DOS Insects*, Chapter 2. The result is that users will have more sense later, the DOSMGR ID field is a trivial buffer, control packet ID 1, see IN\_21F AX=1683h, and instead be 0, as you may have seen in implications for the handling, see "patching DOS" (i.e., by subfunctions 0 and 1 need to be passed down to DOS, with using NODOSMGR WIN 0).

These results may be evidence of what purpose the DOSMGR called API really serves, and what it perhaps is not. The "reality" of DOSMGR resides almost entirely in the magic numbers. This in turn raises the possibility that the real purpose of the DOSMGR call in AP's to the Windows and MS-DOS software is to communicate in a way and not to serve. The interface window drawing as it were. This is a very "poised" life, not one that is worth exploring.

This possibility is underlaid by the fact that some pieces of the interface are not implemented by MS-DOS and that some implementation variations. The current implementation of the interface in MS-DOS 6.0 looks like this, compare Table 1-1.

```

func 0  ret  cx-1, es:bx -> DOS_DS:1022 (DOS var patch table)
func 1  ret  bx:dx, dx:axmagic
func 2  not  handled
func 3  if  (dx & 1) return  cx:5bh (CDS entry size), dx:axmagic
func 4  ret  dx = 0 (trivial)
func 5  get  device driver size from MCB, ret  dx:axmagic

```

The real question is why DOSMGR requires that an invoking DOS 5.0 or higher implement this interface, when DOSMGR also contains tons of hard-coded assumptions about DOS internals. While it is fairly likely that this interface was designed for the sole purpose of making Windows inhospitable to other DOSes, the result is certainly that the interface constitutes yet another "stroke" to Windows' potential, by simply providing very little benefit in return.

At the same time, there is also a positive lesson to be drawn from the DOSMGR-called API. The point has been made several times in this chapter that it is *good* for Microsoft to work DOS into

as the platform for Windows. The impression was perhaps created that this requires intimate knowledge of Windows by MSDOS and of MSDOS by Windows. Certainly that is the impression Microsoft wants to create. But the existence of the DOSMGR callout API (though not its current or future implementation) shows that it would be possible to create a nice, clean, open platform for the DOS of Windows. Were the DOSMGR callout API taken to its logical conclusion, Windows would be a portable, portable, knowledge of DOS internals. All it would require would be an understanding of DOS that is beyond the DOSMGR callout API.

Such a system would not only be good engineering practice, remember modularity, but it would also have a significant advantage. Other DOSes could run Windows and other software on a variety of configurations of DOS. For Microsoft would benefit from such a “Chinese Wall” between Windows and other DOSes. For other DOSes, the benefit would be as with Windows and DOS, only shared without affecting each other. On the other hand, this problem disappears in Chinese DOSes. Windows 4.0 could be sold separately from the underlying operating system.

### Implementing DOSMGR Functions

The variable `CDOS_SIZE` in Figure 113 is used throughout DOSMGR. For example, DOSMGR uses a `copy` option to be removed function shows in Figure 114, that copies the CDOS on the VM specified in `ES:EDI` to another specified in `EDI`. The `SHR` and `CDOS_SIZE` are the CDOS size and VM with the CDOS of the System VM.

Figure 114: Implementation of DOSMGR Copy VM Drive State

```

DOSMGR_Copy_VM_Drive_State proc near
D1E15  pu,shad
D1E16  mov esi,dword ptr [esi+4]    , ESI=source VM, VR [BE4]=base addr
D1E17  add esi,dword ptr CDOS_LIN  , linear address of source VM CDOS
D1E18  mov edi,dword ptr BUFFER   , allocated with _heaplicate
D1E19  mov ecx,dword ptr CDOS_SIZE , number of bytes in CDOS
D1E20  cld
D1E21  rep movsb                    , copy from source CDOS to buffer
D1E22  mov esi,dword ptr BUFFER
D1E23  mov edi,dword ptr [ebx+4]   , EDI=target VM, VR [BE4]=base addr
D1E24  add edi,dword ptr CDOS_LIN , linear address of target VM CDOS
D1E25  mov ecx,dword ptr CDOS_SIZE
D1E26  cld
D1E27  rep movsb                    , copy from buffer to target CDOS
D1E28  popad
D1E29  ret
DOSMGR_Copy_VM_Drive_State endp

```

Such a device driver called DOSMGR Functionally on undocumented DOS. For example, DOSMGR Add Device adds a DOS device driver onto the device chain for a VM. The function goes to the `22h` of `SystemVM.DOS` to find the `NTL` device and the root of the device chain. It then follows the device chain until it reaches the `NTL` of the `NTL` and links in the new device driver. This procedure works exactly like the `DEVICED` program in Chapter 7, although DOSMGR Add Device only works on character device drivers. The Windows `VR6MIMGR` and `SystemVM.DOS` Add Device is a bit more complicated. It MM driver on the DOS device chain unless `SYSTEMINI` includes the line `NOEMMDriver=ON`.

Figure 115 shows the implementation for DOSMGR Add Device. Most of the work here involves getting `SystemVM` getting the root of the device chain, finding the end of the device chain, and so on. The actual work of adding in the driver takes only a few lines (see the label `ADD_DRIVER`).

Figure 115: Implementation of DOSMGR Add Device

```

DOSMGR_Add_Device proc near
00B1A  pushad
00B1B  mov esi,ebx                    , EAX = addr of device header

```

```

00B18    shr esi,4
00B20    shl esi,10h
00B23    mov si,ax
00B26    and si,0FH                ; ESI = seg:ofs ptr to new dev
00B2A    or ebx,ebx                ; EBX = VM handle
00B2C    jnz short GOT_VM
00B2E    VMHCal. Get_Sys_VM Handle , if VM=0, use System VM (g_0ball)
00B34    GOT_VM
00B34    add eax,dword ptr [ebx+4] , VM_CB[4] is addr of VM memory
00B37    mov edi,eax              ; EDI = linear addr of dev
00B39    test byte ptr [edi+53,80h , dev[53] is attrib
00B3D    jz short ERROR          , only character devices supported
00B3F    mov edx,dword ptr LAST_DRV_LIN
00B45    or edx,edx               , already have LAST_DRV_LIN var?
00B47    jnz short NEXT_DRIVER   , if so, walk to end of chain
00B49    mov eax,SYSVARS_LIN
00B4E    or eax,eax              , already have SYSVARS_LIN var?
00B50    jnz short FIND_CHAIN_ROOT , if so, use NUL to find dev chain
00B52    mov eax,reference data , if not, another way to get sysvars?
00B57    shr eax,10h
00B5A    mov SYSVARS_LIN,eax
00B5F    FIND_CHAIN_ROOT
00B5F    add eax,dword ptr [ebx+4] , VM_CB[4]
00B62    mov esi,dword ptr [eax+22h] , in 31h, sysvars[22h] is NUL dev
00B65    NEXT_DRIVER
00B65    movzx ecx,ax            , CX = offset of next driver
00B68    shr eax,10h            , AX = segment of next driver
00B6E    shl eax,4
00B6E    add eax,ecx            , EAX = lin addr of next driver
00B70    mov edi,eax            , EDI = 0 based lin addr
00B72    add eax,dword ptr [ebx+4] , EAX = first lin addr of next drv
00B75    mov eax,[eax]          , get far ptr to next drv
00B77    cmp ax,0FFFFh         , if segment -1, at end of chain
00B7B    jne NEXT_DRIVER       , otherwise, keep going
00B7B    mov dword ptr LAST_DRV_LIN,edx , 0 based lin addr of last drv
00B83    mov LAST_DRV_LIN_VM,eax , first lin addr of last drv
00B88    END_OF_DEV_CHAIN
00B88    add edi,dword ptr [ebx+4] , EDI=1:at lin addr of next drv too
00B8B    mov eax,[edi]         , get far ptr to next drv
00B8B    cmp ax,0FFFFh         , if segment -1, at end of chain
00B91    je short ADD_DRIVER    , if so, can add in driver
00B93    movzx ecx,ax          , still not at end of dev chain?
00B96    shr eax,10h          , try same stuff again
00B99    shl eax,4
00B9C    add eax,ecx
00B9E    mov edx,eax
00BA0    jmp short END_OF_DEV_CHAIN
00BA2    ADD_DRIVER
; EDI holds the linear address of the new driver header
; EDX holds the linear address of the last driver header on the chain
; dword ptr [drv ver header + 0] is a real mode ptr to the next driver
00BA2    mov dword ptr [edi],0FFFFFFFh , make new drv's next = 1 and)
00BA8    mov [edx],esi         , link new drv into chain?
00BA8    cld                   , carry clear = success
00BAB    DONE:
00BAB    popad
00BAC    rsm
00BAD    ERROR
00BAD    stc                   , carry set = failure
00BAE    jmp short DONE
DOSMGR Add_device endp

```

Unlike `DEVICD`, this function does not call the device's initialization function. This behavior is documented in the Disk Access Manager. Also, unlike `DEVICD`, this function adds a device to the end rather than to the front of the device list. As the DDK points out, this means that the service won't help replace an existing device.

The similarities shown in Figure 1-15 could have been set in a number of different ways elsewhere in `DOSMGR`. Given `DOSMGR`'s reliance on the shifting sands of undocumented DOS, it's good that `DOSMGR` contains so much redundancy.

What if the question is whether it is right for one Microsoft product, Windows, to rely so much on undocumented aspects of another Microsoft product, DOS? It is truly a good thing that Microsoft is bringing DOS and Windows closer together and integrating them so tightly, but you also have to remember that Microsoft still sells them separately. If Microsoft uses undocumented DOS calls in its own separate products such as Windows, shouldn't it document these calls? After all, the calls have proven to be useful. Doesn't this imply to document these calls or to provide new, clean, documented interfaces with equivalent functionality, like `IOCH HIDE PIPE` in Windows 3.11? Microsoft gets a gross advantage in bringing out products such as Windows. Is this advantage unfair? Remember, utilities may seem like part of the operating system, but they are one of the most active parts of the PC software market.

Even if Microsoft were not a regulated company that owns 90% of the highways, and that's a case in a number of jurisdictions, now imagine that Microsoft owns hidden features of its highways. This is not to suggest that Microsoft's cars are better than the competition's. The competition can use these hidden features, too, but first they have to somehow find out about them. Many of these hidden features are well known and constitute a kind of tollbooth. But hiding them raises others' cost of using them, and there's always the danger that Microsoft will at some future date remove or change them.

**Patching DOS** Rather than installing the private `DOSMGR` callback API, mostly involves patching DOS. We really do mean patching here, not hooking an interrupt, though `DOSMGR` and other pieces of Windows do plenty of that too. The API writes over bits of DOS with `REP MOVSB` instructions.

In MS-DOS 5.0 and higher, `MSDOS.SYS` responds to an `INT 2Fh` `AX=1607h` `BA=15h` `CA=0` call from `DOSMGR` by passing back in `ES:BX` a table with pointers to six DOS internal variables: `SAVEDEV`, `SAVEERR`, `NDOS_USER_ID`, `CRITPATCH`, and `UMB_HEAD`. These names come from the Microsoft internal focus on the `DOSMGR` API; presumably these names correspond to the actual variable names in a DOS source code, one of our tech reviewers writes, "they do!" The patch locations aren't in the source segment as the patch table itself, so only offsets are provided. Figure 1-16 shows a hex dump of the patch table in MS-DOS 5.0 and 6.0.

**Figure 1-16. DOSMGR Patch Table in MS-DOS 5.0 and 6.0**

```
0116 1022 1021
0116 1022 0005 05FC 05EA 0321 033E 0315 008C
```

```
Patch table at 0116.1022
DOS version 5.00
SAVEDEV: 05ECh
SAVEERR: 05EAh
INDOS: 0321h
USER_ID: 033Eh
CRITPATCH 0315h
UMB_HEAD 008Ch
```

What does `DOSMGR` do with these pointers? Here again we see the half-hearted nature of the `DOSMGR` callback API. Basically, `USER_ID` and `UMB_HEAD` are the only variables of any great importance.

`USER_ID` is which DOS keeps the machine ID. As shown in Figure 1-17, `DOSMGR`'s handler for the VM Critical Exit message, which VMM sends out whenever a new virtual machine is starting,

up, smacks the new VM's ID number into this location in DOS. DOS in turn uses `USER_ID` to stamp the SFT entries for specifics with the VM ID (see `HEADC` in Chapter 8). We'll explore the cryptic variable names in a moment.

### Figure 1-17: DOSMGR Patches the VM ID into DOS's USER ID

```

; EBX contains the VM handle, a pointer to the VM control block (VM CB)
mov eax, dword ptr [ebx+CB_VRID] ; get VM ID # from VM CB
mov edx, dword ptr LIN_PATCH_USER_ID ; linear addr of USER_ID
add edx, dword ptr [ebx+CB_high_linear] ; base address for VM's memory
mov [edx], ax ; smack in the VM ID #

```

In Figure 1-17, DOSMGR is patching the VM ID into the new VM's high end of DOS's data segment. As we will see later, DOS normally sets the machine ID to 0's to show DOS knows that Windows Enhanced mode is running; it skips this operation. DOSMGR gets the address of `USER_ID` in the DOS data segment using the `VM` control patch table, pointed to at 16. When a program opens a file, DOS passes the program's PSP and the value of `USER_ID` into the file's SFT entry. Under Windows Enhanced mode, the same set of call SFT entries not only the owner's PSP but also as VM ID. This is used, for better or for ill, by the combination of the handle and PSP. DOS notes Windows Enhanced mode uses the combination of the handle, PSP, and VM ID. DOS needs to do this to do "logical" PSP's, different in DOS because not possible as appears as the same file, since `SYSTEM` seems as much as `USER`. DOSMGR's `USER_ID` PSPs in turn also address this problem. The whole point of all this procedure is to keep track of multiple DOS files, containing on top of a single copy of DOS.

The code fragment in the step show that DOSMGR uses a variable `LINE_PATCH_USER_ID`, which is the 32-bit linear address of the `USER_ID` patch location. As `USER_ID` is 32-bit program, so they can't really vary, even by machine, from the address. In program patching, we can make it set such as `MOV EBX, VM`. Show the code that DOSMGR sets it, creates the `LINE_PATCH_USER_ID` pointer, and requires the patch table. The assembly code is shown through the step show in Figure 1-18. The pseudocode in Figure 1-18 may also be parallel some of the code shown in earlier figures.

### Figure 1-18: DOSMGR Forms a Linear Address to DOS's USER ID Variable

- Use the `MMIO` function `ExecFid` to generate a 32-bit `MMIO` handle, which returns a pointer to the DOS `SVARS` also in `USER_ID`. `ExecFid` makes these variables in the `C:\DOS` and `C:\WINDOWS` subdirectories, a control structure, accessed off by `EBP` (exists, see Figure 1-8).
- Use the `ExecFid` function `ExecFid` to generate a 32-bit `MMIO` handle, which returns a pointer to the DOS `SVARS` also in `USER_ID`.
- Set the `DOS_VN_CB` word in `VM` to a 32-bit `LINE_PATCH_USER_ID` variable, 32-bit linear address of the DOS data segment.
- Use `ExecFid` to generate a 32-bit `MMIO` handle `VM` and `LINE_PATCH_USER_ID` the DOSMGR `VM` call out. DOSMGR uses a pointer to the patch table in `USER_ID`.
- Set the `C:\DOS` `USER_ID` word in `VM` to a 32-bit `LINE_PATCH_USER_ID` variable, 32-bit linear address of the patch table.
- Call `MMIO` function `GetC` or `VM` handle to get the VM handle, which is a 32-bit linear pointer to the VM control block `VM CB`. Offset 4 in the `VM CB` is `CB_High_Linear`, which is the 32-bit linear address of the VM's memory in the Windows linear address space.
- Add a `VM CB_High_Linear` near on to the 32-bit address of the patch table to get the 32-bit linear address of the patch table within the entire Windows address space. Call this `LINE_PATCH_USER_ID`.
- The remaining steps are easier to show as code.

```

mov esi, LIN_PATCH_TAB
mov ax, WORD PTR [esi+8] ; 4th word is USER_ID
mov PATCH_USER_ID, ax
movzx eax, word ptr PATCH_USER_ID ; zero-extend into EAX
add eax, dword ptr DOS_05_LIN ; add in linear addr of DOS_05
mov LIN_PATCH_USER_ID, eax

```

On the subject of conventions, MSDOS.SYS simply has an INT 2fh handler that looks out for AX=1605h (i.e., SUBDISMGR) in Fastlog 1's patch 1's patch 1's INT 2fh handler ahead of this one in MS-DOS.SYS (DOS.SYS's INT 2fh AX=1605h, if with BX=15h, it checks the subfunctions in AX). For subfunction 0, it puts the address of a user-based table (i.e., one shown in Figure 4-16) into EBX.

From 0325h, the patch table (with the user ID's called CRITPATCH) is torn in critical sections (DOS patch 1, DOS S.O., one) and repaired (i.e., another, at different locations in the DOS code). The patch is stored on the C:\C:\C:\C:\C:\C:\C:\C:\C:\C:\C:\C:\C:\C:\C:\C:\C:\C:\C\COPY\ directory. Starting with DOS S.O., there is one central routine that all the DOS critical-section code checks. This is a table in the DOS data segment, which is crucial for DOS in RAM because you can't patch it when you're in operation (DOSMGR does these patches). Its location (to form an critical sections) is the CRITPATCH parameter in Chapter 9, and the appendix's earlier INT 2fh handler 800h, for instance. What's so special about it and backs out all these patches when it exits? Well, usually," writes one of our tech reviewers.

It creates LMB (LMA) which Windows 3.0 checks for one's the address of a word where DOS would be programmed with the last meta-handler into conventional memory, that is the base of the upper 640K block LMB (which is location is given instance 1 for each VM).

## DOS Knows About Windows

As for the latter, when Windows starts up, it issues a INT 2fh call with AX=1605h, when Windows exits it issues a INT 2fh call with AX=1606h. MS-DOS S.O. and 6.0 hook these calls to maintain (for 1.31a and mode 0) a IN\_WINDOWS flag. When Microsoft documents the 1605h and 1606h functions, it says (1605h): "This is a deprecated call and the documentation states that this is to notify device manufacturers. SR for Windows 3.0 setting. The documentation says nothing about the important feature of any of them that MS-DOS itself responds to these Windows calls."

MS-DOS sees the INT 2fh AX=1605h Windows initialization broadcast for a number of purposes, for example, Windows 3.11 (in real mode) starting up, DOS S.O. and 6.0 tell Windows to read the WIN32 flag (i.e., these patches passed to the Windows 3.0MMGR). MS-DOS also uses the message to read Windows score (i.e., system data). Windows score provides a lot of different ways to specify instance data. Again, Microsoft documents this interface in the DDK, merely failing to note that MS-DOS itself uses the interface.

In DOS S.O. and the Windows 3.0 MMGR, one can not only tell Windows about instance data and VxDs (i.e., 3.11's VxDs), but also tell it what indicates whether Windows Enhanced mode is currently running. DOS S.O. does this in a way or the other (i.e., a word's value) and, for when Windows Enhanced mode is running, MS-DOS wants to be able to do a network-wide cleanup. If you think about it, multiple DOS instances Enhanced mode are very much like multiple machines in a network. As noted earlier, DOS S.O. has a way to ID number (i.e., the INT 2fh dispatch loop normally sets this ID to 0 and it can be overwritten, see Chapter 6). However, if DOS sets the IN\_WIN32 flag, it doesn't set a unique ID number, to be instead, issues the current DOS box's Virtual Machine ID number (see INT 2fh AX=1680h). As we saw earlier, DOSMGR patches this VM ID number right into DOS, which then copies it into every SFT entry.

DOS.SYS also sets the IN\_WIN32 flag and, if Windows Enhanced mode is running, DOS.SYS will do a VxD check (i.e., instances) make a call to DOSMGR. In addition to the INT 2fh AX=1607h initialization call that VxDs can make, VxDs can also provide APIs. These APIs are available even to software that the DOS.SYS and DOS 6.0 (DEINPAC.BIN) was loaded before Windows. Basically,

VxD APIs provide a way to use VxD code (via normal I/O) but DOS code (to give 32-bit protected mode VMM and VxD code execution from the chipset, see "On Architecture and How Anything Works" by Virtual Device Drivers for Windows™, *Microsoft Systems Journal*, October 1992, Volume 1, No. 2). For AX=1684h with the VxD ID, such as 15h for DOSMGR, WinX removes a VxD API entry point. The DDK documents this INE 21h mechanism. What Microsoft does not cover is that the actual INE 21h AX=1684h API is provided by DOSMGR and other VxDs but into Windows. As with INE 21h AX=1685h and AX=1687h, the DDK documents only the general mechanism, not the way that Windows and DOS actually use the mechanism.

The undocumented DOSMGR API has some functions: DOSMGR\_DESPAWN calls setstate, not 1; this is a function in generic VMM4.dll. When New Game's function is installed, callback function that updates the state. Also Adjust Execution Time affects some non-ideal software's execution.

DOS also uses the word set state as a method to indicate whether to generate NtVxD AX=8001h and AX=8002h error sections (see "state"). DOS skips these calls as they are needed to mark unimportant locations of code in a single-tasking operating system. The following flag is exactly adequate (see Chapter 9). However, Enhanced mode effectively turns DOS into a genuine preemptively multitasking operating system.

Again, while it is a bit surprising to hear that DOS removes a flag that Windows Enhanced mode is turning, it is important to realize that DOS makes this change using a documented function call. The Microsoft DDK documentation should not state clearly that MS-DOS user books this call.

### **DOSMGR and the SDA**

Back in INE 21h output on 16-bit, you can see that DOSMGR not only calls INE 21h AH=52h, but also calls INE 21h AX=8100h. As noted in the appendix, this window move function results in a pointer to the DOS Swap area (Data Area SDA). The SDA is a pointer into the structure in the DOS data segment that contains the current DOS state, including the current Program Segment Prefix (PSF), the current disk transfer byte (DTB), the DOS flag, the current drive, the three DOS stacks, and so on.

DOSMGR calls INE 21h AX=8100h because it needs to mark the SDA as instance data. Whereas DOSMGR has to go through various controls to figure out the size of other DOS internal structures that must be assumed to exist to find out the size of the SDA. As an example, appendix INE 21h AX=8100h needs a pointer to the SDA, also contains a number of bytes that a program must always swap, and the number for a program need only swap if the INDOS flag is enabled.

DOSMGR passes these methods directly to the VMM. Although the PC is not which, it is, essentially, accepts instance-specific INDOS flag and MWRV flag. Given that the SDA itself is undocumented, the flag is not really the INDOS flag, it is a word in swap area storage that the DDK does invents these flags especially because VMM passes to some third party way types of the program's system calls by using the document for DOS and Windows. The Windows documentation frequently references over pieces of DOS knowledge that is essential to understanding.

### **DOSMGR and the Indos Flag**

As a good example, consider the DOSMGR\_Get\_IndosPtr function. Documented as the Windows DDK. According to the DDK VDDA, this function returns the linear address of the MS-DOS INDOS and Error Message inspection of the code in DOSMGR (see Figure 1-19) confirms that this is indeed exactly what the function does.

**Figure 1-19: Implementation of DOSMGR\_Get\_IndosPtr**

```
DOSMGR_Get_IndosPtr proc    near
01707  mov  eax,INDOS_PTR
0170C  ret
```

```
DOSMGR Get_IndosPtr endp
```

```

; In initialization
05A7A mov [ebp,Client_Ax],3600h ; set up client_reg struct for call
05A80 mov eax,21h ; call INT 21h AH 34h via Exec_Int
05A85 VMRCall Exec_Int ; (Get [DOS Flag Address])
05A88 movzx eax, [ebp,Client_E5]
05A8F shl eax,4 ; ES <= 4
05A92 movzx ebx, [ebp,Client_BK] ; mov with zero-extend
05A96 dec ebx ; back up one to get ptr to CritErr
05A98 add eax,ebx ; create 32-bit linear address
05A99 mov INDOS_PTR,eax ; actually, both Indos and CritErr!

```

DOSMGR backs up the pointer returned from INT 21h AH 34h to get both the IndOS and CritErr flags into one convenient word-sized chunk, which can be used to test both IndOS and CritErr in a single operation.

While it's so obvious, it's not documented that INT 21h AH 34h and the IndOS flag are documented until the *MS-DOS Programmer's Reference* that Microsoft issued at the time it released DOS 5.0.

Not only does the book show in Listing 3.15 writes a flag before that bit, more important, Microsoft's MS-DOS documentation is still missing a critical detail. It fails to note that, in DOS 3.0 and later, the value stored in the IndOS flag is the *critical error flag*. Microsoft probably doesn't document this because, in some versions of DOS, the critical error flag was set elsewhere (see the appendix "Yet Another Table" that shows that the undocumented DOSMGR Get\_IndosPtr function relies on this undocumented aspect of DOS). Actually, as this book was going to press, we found that Microsoft's DOS programmer's reference *did* refer to the obscure place to the fact that the critical error flag includes the IndOS flag. In the words of Emily Latack: "Never mind!"

## SYSTEM.INI Settings and Undocumented DOS

Not only does Microsoft have programmer's documentation for Windows that relies on things that are undocumented in MS-DOS, but it also has *end user* documentation for Windows that relies on things that are part of undocumented DOS.

One Windows feature—some way to DOSify interrupts—clears even from a non-programmer's or end user's perspective, because the Windows SYSTEM.INI configuration file can contain settings and relate to DOS operations and structures that are, or were recently were, undocumented. The SYSTEM.INI file that is used for Windows describes those settings, as does the *Windows Resource Kit* and *Windows for Workgroup Resource Kit* that Microsoft publishes.

```

[star.dard]
Int28Filter=10
[386enh]
IndosPtr=regNo
Int28Critical=True
PSPIncrement=2
ReflectDosInt2A=False
UniqueDOSPSPT=True

```

For example, the SYSTEM.INI file, which is normally assumed at end users, not programmers, and which comes with the retail version of Windows, contains the following description of ReflectDosInt2A:

```

ReflectDosInt2A=Boolean
Default=False

```

Property: Indicates whether Windows should consume or reflect DOS INT 2A signals.

The default means Windows will consume these signals and therefore run more efficiently.

Fixed: This setting is not available if you are running the more resource software that relies on detecting INT2A messages.



Microsoft's programmer documentation, even the latest programmer's reference for DOS 6.0 has no mention of INI 2M's second stage (and, incidentally, critical sections). While apparently indicating that DOS programmers don't need to know about INI 2M, Microsoft clearly still needs to say something about it to Windows end users.

In Windows 95, both Windows 95MSMR and its successor VxDs check these SYSTEMINI settings. The ReflectDOSIni2M setting, along with an undocumented ModifyDOSIni2M setting, is also checked by 95MSMR, which contains an INI 2M handle. This is another window for Windows Hook VxDs (or C++ functions) to install traps for INI 2M's various appropriate functions provided by the Windows VMM. For example, the 95MSMR INI 2M handles its Register Critical Section on to implement INI 2M's All 80h End Critical Section on to implement INI 2M's All 81h and All-82h, and Release Critical Section on to implement INI 2M's All 84h. If ReflectDOSIni2M is set, then 95MSMR also passes down "reflects" INI 2M calls to any INI 2M handlers that were installed before Windows.

### **KRNL386 Grows the SEF**

Well, that's enough for 95MSMR. Popping back up to the INTRNPs results in Figure 1.6, as we see that KRNL386 INI 2M in Windows kernel's initial configuration, and it recalls our old friend, INT 21h function 52h to get a pointer to SysVars.

KRNL386 uses SysVars to get a pointer to the System File Table chain. Windows takes a subtle but a use of files that normally would be used to keep a list of the open handles, so KRNL386 tries to increase a number of SEF entries by using the same undocumented DOS APIs that entries like Query, Query Files, and I/O Process's MIDLINOS. Actually, 95MSMR also can grow the SEF by depending on the value of the SYSTEMINI setting of SYSTEMINI and whether SEFRES is loaded.

KRNL386 expands the number of SEF entries during its boot sequence, through BootStrip, BootStrip, an internal KRNL386 function, calls another internal function, GrowSEF and May do the dirty work. Chapter 10 of Matt Pietrek's book *Windows 95: The Implications of the Windows Operating System* (Etc., Inc., New York, 1995), "BootStrip" and "GrowSEF and May" provides a listing of INI 21h API's 52h, GrowSEF, May, grows the DOS 21h's internal Mode structure, DPMI, to get a protected mode section to SEF entries of the SEFs. After determining how many entries the SEF will create, GrowSEF and May multiplies this number by 4, EntrySize, and calls the Windows GlobalData, etc., function to allocate the new SEF blocks (entries) from memory. It constructs them into the new SEF blocks into the end of the existing SEF chain. And, MaySEF finishes.

When it's time to use the value of a SEF entry, come from the depends on the value of SEF entries 40h bytes of DOS 3.0 48h bytes of DOS 3.1 33h and 48h bytes of DOS 4.0 and 4.0. But 95MSMR does not rely on its clear boot-time knowledge of SEF 386's internal InIniVxD's function calls yet another internal function, GetEntry. You'll remember, now KRNL386 figures out the size of an SEF entry. It opens a CON file handle and scans the first 512 bytes of the file for the string "CON", just as we saw 95MSMR doing earlier. Patrick's *Windows Internals*, which supplies pseudo-code for InIniVxD, expands call to New with the constant "Brett" as "Null" and

### **KRNL386 and the PSP**

Finally, KRNL386 calls undocumented INI 21h API's 55h to create a PSP function.

As explained in Chapter 7, every process in DOS has a PSP; the DOS EXEC function, INI 21h API's 4Bh, automatically calls INI 21h API's 55h to create a PSP. The PSP contains process information from the most important of which is a pointer to the PSP's file table. All specifies in DOS are associated with a particular PSP. For KRNL386's paragraph address also acts as an identifier for the process. For example, DOS Memory Control Blocks are generally owned by some PSP. When the process exits, any MCs marked with the block's PSP are freed, unless, of course, the process exits via the DOS terminate and stay (TSR) function.

in real mode DOS. Why is the Windows kernel creating PNTs? Because Windows isn't actually all that different from DOS. Windows is a protected-mode DOS extender. Every process in Windows *always* a PNT. Why? A process in Windows opens a file, its file handles are associated with you, guessed it: the PNT. When a Windows program allocates memory via GlobalAlloc, the block of memory is except for GEMM (SHARE allocations) associated with the PNT.

Now, PNTs don't even exist until the program is in a stable state of a Windows process. Because Windows runs on a more transparent DOS. For example, when running Windows program has a process of 100, and there's no scheduling. Taking the program needs 2 links to nodes, and Windows implements a message-passing, there's no room for any of this in the PNT so Windows uses a different structure, the Task Database (TD). For simplicity, Windows uses its PNT as a process database (PD). Actually, one of our tech reviewers for *Windows 95* called PNTs as called PDBs in the DOS source code. As both *Undocumented Windows* and *Windows Internals* explain, the TD is an SP-11M structure together. You can think of the TDs as an extended PNT.

First, the KRNLS00Data is a new task structure of data structures for the task, called task PNT. This is a genuine DOS PNT, allocate for conventional memory, where DOS can get at it. This is why Windows can't grant permissions, low-DOS memory, an insufficient amount of conventional memory, a Windows program can't running, every other is otherwise priority of memory. After all, a Windows program opens a file, DOS will need to get a PNT to store the file handles, and PNTs may be used to low memory, order for DOS to read them. An important side point is that to respond to PDs, they are really associated with whatever task was running when the file was opened. All real tasks and low-mode PNTs.

So, you think you can Software, WinWord, Excel, Ami Pro, or any other Windows program, the kernel, `NTL`, `21h` function `58h`. As explained in *Windows Internals*, this call invokes the internal kernel function, `CreateTask` and `FilePNTs`. Why then, didn't we see a whole mess of `NTL 21h` API's in our task and process tables? Because, the first objects passed down to real mode DOS — Windows is a real mode, KRNLS00 handles these calls. You can still see the protected mode `NTL 21h` API's of `NTL` in `Windows`, however, it's a non-Windows interrupt trapping program, called `WINPNT`. I speak Windows, *Undocumented Windows*.

There are many more aspects to Windows and the PNT, but we concentrate on what I have room to get into here. For example, what is the appendix entry for the PNT (see `NTL 21h`, function `26h`)? Windows, see the control table in the low memory purposes. Division of object `42h` forms a PDB. In a real mode, `40h` indicates whether a program running under Windows is an "old" DOS program or a protected-mode Windows program. Chapter 4 discusses the Windows PNT in more detail.

## Undocumented DOS and the Utilities Wars

For a real look out from Windows to look at other systems software for the PC. How much do disk compression and disk utility programs, caches, indexes, managers, DOS extenders, networks, and so forth on undocumented DOS call?

DOS and some DOS utilities such as SUBST, JOIN, SHARE, CBINT, AMM, and the Local use index, `INDEX`, `INDEX`. You can check them out for yourself, without any of these programs under `NTL 5` using a `NDDIR`, `SCR` from `1.2` of these programs and `NTL 21h` function `52h` and `www.nv.com`. It is by the way, `NTL 5` was the DOS's own internal use that function `52h` was created in the first place.

There are a lot of other utilities, but every operating system must, by its very nature, have multiple users. One of the key purposes of an operating system is to shield users and programmers from computers, and the OS utilities such as `NTL 5` do that. It is vendor documented every last detail, only, users and every suggestion, but in the PC marketplace, utilities are almost the application. So, with the fierce competition, and some of the largest sales come from utilities. It

is not clear why they are so incredibly popular with consumers, but the fact is that XMS and EMS utilities take a surprisingly important place in the world of top-selling PC software. In a randomly selected issue of *PC Magazine* (March 1993), the top ten retail software packages included two of these: QEMM 386 v.0, Quarterdeck, and Stacker 3.0, Star Direct ones.

The interesting thing is that all of these top-selling PC utilities rely on undocumented DOS functions and data structures. Many of the programs that make up QEMM, including GUEFNOCOM, I4BS.COM, I1E5.COM, I4S1DRIV.COM, I4ADH.COM, and I4I1E1.COM, call INT 21h functions 52h. In Stacker, the programs I4NDVH.COM and I4CHECKI4S1ACKI4R.COM, SWAPM4P.COM, and S4REI4I4I4E1 call this most popular of undocumented DOS functions in order, for example, to find the DOS Disk Parameter Block (DPB) entries. Writing a competitive DOS utility without knowledge of undocumented DOS functions and data structures is impossible.

If utilities are a hot area of competition in the PC market, what is yet to compete in? The market depends on being informed on about DOS that Microsoft refuses to disclose. If then Microsoft's refusal indicates a major problem. As a result, if MS-DOS incorporates innovations, such as memory management and disk compression that only require support to products from other vendors but which probably belong to the operating system, many third-party or independent QEMM and Stacker must increasingly compete with DOS itself. Microsoft has a very large legal staff producing software for the market and is control over the operating system documentation on which the playing field even more uneven.

On the other hand, lack of documentation hasn't yet kept any one out of this market. It is, as we've noted, quite competitive. While Microsoft does document everything about DOS that it should, programmers can pick up this information through reverse engineering or through warez, such as this book. For all we know, this comes from programmers, Microsoft use find out about DOS internals.

The way that utilities use undocumented DOS, as well as the way Windows uses it, is a important question. Is it right for Microsoft not to document functions that are used for competing in the PC utility market. While the more the better to this question, it's important to note how the use of As we did with Windows, we can poke around in the different operating systems, some with DOS and some what use their make of functions that Microsoft doesn't document. In other examples of this, our mention about undocumenting DOS and PC utilities, we first look at SmartDrive, DoubleSpace, and MMIOs that we'll examine some of. In Chapter 3, third party utilities that are being developed for a MS-DOS 5.0 and 6.0, MIRROR, DIRECTRAK, and EMASoft, are also.

### Undocumented SmartDrive

Long referred to as DiskDance because of its poor performance, Microsoft's SmartDrive disk cache has seen many improvements in the past, where it's now a viable alternative to commercial products such as M. M. K. Cache. In DOS 6.0, you can use SMARTDRV.EXE or an alias from the DOS command line, in contrast to the older SMARTDRVSYS which you could only load as a device driver from CONFIG.SYS. On the same hand, the DOS 6.0 version of SMARTDRV.EXE also has a much aggressive anti-theft behavior, which is crucial to data owners in the world of experienced users.

SmartDrive calls INT 21h function 52h, but it seems like every program we look at uses this function, so this is no longer very interesting. What is more significant is the fact that SmartDrive itself has an undocumented API. For years, Windows programmers knew that Windows was able to grab memory back from SmartDrive temporarily. But how it did this, and how non-Microsoft programs could perform the same trick as Windows, was shrouded in mystery and confusion. Microsoft never documented this in detail, but after reverse engineering SmartDrive, Geoff's happy revealed all in a January 1992 *Dr. Dobbs' Journal* article, "Untangling SmartDrive." Chapter 5 also shows

447 programs, and manipulate the SmartDrive cache using an IOCH write (INT 21h AX=4407h) and an IOCH read (SmartDrive using IOCH read INT 21h AX=4402h).

SmartDrive 4.0 and 4.10 do not converge to the IOCH interface described by Chappel. SmartDrive 4.0 uses INT 21h AX=4408h instead of this interface, sometimes referred to as the BARE interface since it resembles BARE's spinwait timer, is also supported by Vx Cache 8.0 (see the Appendix). DOS 6.0 comes with the SMARTEMDS utility that uses this interface to measure various aspects of SmartDrive's performance. Jeff Prouse's MS-DOS Q&A column in *Microsoft System Journal* (September 1992) repeats this mistake. Finally, but Microsoft has not officially documented it, an *msn* case file showed error messages after the presence of SmartDrive 4.0 (not 4.1) to shrink its cache temporarily. It appears that to shrink the SmartDrive 4 cache, a program must pretend to be Windows by issuing the INT 21h AX=1605h and AX=1606h initialization and termination calls.

### Undocumented DoubleSpace

Veritas Software, Inc. has an impression of it provided by Microsoft in MS-DOS 6.0, is based on Veritas software. DoubleDisk. Its direct incorporation into DOS makes DoubleSpace formidable competitor. DoubleDisk is such as the original disk compression utility. Max Information's Stackler. It is well documented that Microsoft paid Veritas no money for DoubleDisk, offering instead the mere opportunity to market some products to MS-DOS 6.0. If true, this rumor sheds an interesting light on Microsoft's decision to support and incorporate 4 technology from smaller software vendors.

DoubleSpace's early integration into MS-DOS 6.0 is MS-DOS 6.0 a new system file, DBLSPACE.BIN, runs IO.SYS or 1EMBEDDED.SYS. IO.SYS loads DBLSPACE.BIN before processing COMMAND.SYS. DoubleSpace's integration is a way that use the excellent Stackler product's services in a DOS 6.0, where Stackler products use DBLSPACE.BIN, thus sharing in same product benefits as DoubleSpace, in essence, DOS 6.0 has created a new "preload" API. That is, DoubleSpace's system file and DOS kernel are not otherwise largely unchanged from DOS 5.0 to 6.0. It is consistent with the fact that IO.SYS and FORMAT.SYS command now copy DBLSPACE.BIN along with IO.SYS in MS-DOS.SYS. DBLSPACE.BIN has two components: a disk space manager and a compress/decompress "engine" (see Chapter 8).

Microsoft's early Microsoft makes DoubleSpace's incorporation into the DOS kernel its major claim to MS-DOS 6.0's "new Stackler claim. Microsoft tells its DOS licensees such as IBM that DoubleSpace is not part of MS-DOS. In this claim arrangement, the retail DOS 6.0 with DoubleSpace is sold to Veritas, and so IBM does not get DoubleSpace. So much for Microsoft's claim that DoubleSpace's early Stackler's integration into MS-DOS. When it suits Microsoft's purposes, it holds DoubleSpace is which separate from MS-DOS.

Microsoft's dealings seemed some important lessons, because it is making available proper products. The Microsoft on DoubleSpace. The "DoubleSpace System API Specification" documents the INT 21h AX=44xx interface provided by DoubleSpace. Early versions of this specification (see Appendix 7 and Appendix 25) this API is reserved for SmartDrive. However, after completion of development, Microsoft decided that the final version of the specification to document these calls since SmartDrive's use of and not because it is its points with DoubleSpace. These functions are useful for non-Microsoft disk caches.

Microsoft's "DoubleSpace Compressed Volume File Overview" describes the actual disk format of CVFS (see Chapter 8). The "Microsoft Research Compression Interface (MRC) Specification" documents early interface MRC I protocols, which uses INT 21h AX=4412. The MRC I specification of 2000 documents INT 21h AX=4400 for compression hardware.

Finally, DoubleSpace. MRC I documents are readily available on CompuServe, on the Microsoft Developer News, k:MSDN\_CD-ROM, and in the MS-DOS 6.0 version of the *MS-DOS Programmers Reference*. I wish contrasting the abundance of information Microsoft provides for DoubleSpace with the difficulty of finding program's documentation for Stackler.



I would like to implementing some important undocumented interfaces. DoubleSpace itself relies, in a surprising way, on undocumented DOS calls—given its competition with Stacker and SuperStor, this is surprising. Of course, as noted above, Stacker also uses undocumented DOS, and presumably SuperStor does, too. And, as also noted above, Stack's programming documentation for Stacker is nowhere as readable as either the Microsoft's documentation for DoubleSpace, but isn't there a difference between a third-party vendor using undocumented DOS calls and Microsoft using them? And isn't there a difference between a third-party vendor being to document its programming interfaces and Microsoft doing so? I believe that both have on occasion and that, even if a trade practice is legal in the hands of a software company, the same practice can be legal in the hands of a non-polit.

DoubleSpace—DOS SPACELIB with Source9 shows that DoubleSpace calls INT 21h function 52h to locate other files. So what does get a pointer to the DDT's chain and offset 16h to get to the CDS entry? DoubleSpace, in turn, makes a call to the undocumented DOSMGR.VXD in Windows.

INT DOSMGR.VXD—the SPACELIB tries to access parts of what are the high memory area (HMA) and the 80386 extended memory. How does it do that? Programmers have frequently noted that there must be parts of the HMA that DOS doesn't use at, for example, the BUFFERS-setting is small. How does DOS SPACELIB get at these areas? It uses two undocumented calls: INT 21h AX=4401h (QJ) and INT 21h AX=4402h (VJ) in the HMA space (see the appendix). These functions were well-documented in a code sample in *Microsoft Systems Journal*, March 1993, in Jeff Prosser's MS-DOS Q&A column; it is depressing to see a comment in the state of DOS that the ability to access an extra 67K bytes of what's called the HMA is a useful feature, so this is another interface that Microsoft should formally document.

### Undocumented EMM386.EXE

Compaq appears to have originally written Microsoft's memory manager EMM386. EMM386 can protect dual-bank memory managers such as QEMM and 386Max.

EMM386 uses a lot of our knowledge of Intel's 8086 protected mode and virtual 8086 V86 mode. In our knowledge of undocumented DOS, this includes undocumented aspects of the Intel architecture, such as the 286 and 386 L0ADMM instructions, which EMM386 EXE special cases in its most obscure handler. For more on L0ADMM, see the superb article by Robert Collins, "The L0ADMM Instructions," which appeared in *THE OPERATOR* (October 1992).

EMM386.EXE is similar to Windows's virtual device driver VxD. What we normally think of as the DOS driver, that is, the real-time stub, is a so-called VxD. The stub of the file looks like a normal DOS exec in the V-Memory base code library article "hooking a VxD to a VxD" Q74516, describes how subroutines that pass extended 48-bit MMR, DRX, EX, and IN, LINK, INF use the same format.

INT VxD stub of EMM386's called LoadH, whose stub with Source9 shows that, in addition to hooking the DOS Set Upper Memory Limit function, INT 21h AX=5803h, and hooking VxD functions such as VxDInit and VxDExit, it calls the V86MM and DOSMGR Instance Device, the LoadH VxD also makes a call to V86MMGR. This patching is evident when you do a strings list of EMM386.EXE, or the LoadH device, such as 386MAX.VXD.

I would like to cover how a LoadH stub is an example of how a DOS provided utility such as EMM386 goes about patching Windows. According to the DDK, one of the functions VMM provides—hook Device Service. This function allows any VxD to intercept any function that VMM or another VxD provides. This is an interesting aspect of the far-reaching capabilities that VMM functions offer. Hook Device Service takes a return number, for example, Hook Device Service itself is function 020090h, and is similar to the very interesting function Hook Device Service returns the previous handler for the return, a fact which DDK doesn't explain very well. The intercepting function typically sets the pointer to the previous handler to "chain" this is how LoadH hooks the DOSMGR Instance Device call, for example, so far so good.

Now, a VxD that wishes to *pass* VMM and VxD functions (rather than *hook* them) abuses Hook Device Service. The VxD passes in an invalid function pointer (i.e., gives back the previous handler then calls Hook Device Service a second time with the function pointer to the previous handler. Nothing has been changed, except the VxD now has the address of the code it wishes to patch. Generally, the VxD compares some bytes at that address to see that the code matches what it expects, then it smacks in new code. The original position of EMM386 uses this technique to patch the V86MMGR Set Mapping Interrupt V86MMGR User VMS VMS Events functions. Under certain circumstances (i.e., WIN320.386 VxD included with DOS) uses the same Hook Device Service trick as EMM386 to patch the V86MMGR VxD in Windows 3.0.

That an MS DOS reality patches Windows is the most extreme case one can imagine of the special relationship between these two products, but it is not unique. 386Mag and QEMM also include code that so they Microsoft shares information with at least some of its competitors. Note a version of EMM386 in DR DOS provides a `LoadByVxD` for Novell's `nov` that is also received the Windows 3.1 version of the `LoadByVxD` code out of the revised `QEMM` Importing specification and so they had to pay DFL their source. Meanwhile, most other memory managers I've seen do not do anything getting these materials. The same thing seems to have happened with the VMS 3.0 specification. This is a good example of how the problems with Mac software, such patches, undocumented interfaces, but *no* *comprehensive* documentation, which the company gives out to its customer competitors on a case by case basis, is successful.

This seems to be the case with EMM386. For example, EMM386 cooperates with V86MMGR using an undocumented interface, sometimes called the "Windows 386 Paging Import" specification and sometimes called the "Central EMM Import" specification. While Microsoft has not documented this specification, it does appear to have shared this information with some other memory management vendors. It is hard to imagine anyone else who would want this information anyway. One tech reviewer writes, "In fact, I cannot imagine people wanting this information to help them take over from a memory manager." Shresh's).

In addition, EMM386 provides "small" API through `INT 67`, `API V86`, `INT 68`, `API V86`, `INT 69`, `API V86` on the disk accompanying this book.

With all the undocumented things EMM386 does, it is not clear that here, Microsoft has at least some sort of strategic agreement with the other memory management vendors. No, perhaps is a more problem rather than a strategic problem. Also, it is not clear how any of EMM386's tricks give it a competitive advantage. EMM386's `nov` file is more conservative than 386Mag and QEMM using fewer tricks, of course. EMM386's competitive advantage comes entirely from a better handled tree with the operating system, not from using or abusing unique I/O tricks. Finally, it is worth noting that the Microsoft Windows documentation contains a fairly clear statement of the "discrepancy and discontinuity" of EMM386, compared to other third party memory managers.

EMM386 is discussed in more detail in *Goal 4* Chapter's *1995 Introduction*.

## Microsoft Anti Virus

Starting with MS DOS 5.0, Microsoft began incorporating formerly third party utilities into the operating system. DOS 5.0 included the MIRROR, FORMATS, and UNDELETE utilities from Central Point Software, the makers of PC Tools, and competes with the Norton Utilities (and Symantec). DOS 6.0 incorporates more utilities from Central Point, plus Security with RAV, which is similar to SpeedDisk and Norton Utilities. The Central Point utilities in DOS 6.0 include many Windows executables and DLLs, apparently a PC Central contract received by executive's tools to Microsoft was to be used from Microsoft, allowing Central Point to see the "look and feel" of the DOS/HEIF. See *Windows*, June 12, 1991, "Some deal".

We could spend days just looking at the different Central Point utilities included with MS-DOS 6.0. The inclusion in DOS of software that requires Windows is itself interesting, but it is less interesting

... these utilities. Microsoft Anti-Virus (MSAV) In a useful roundup of the utilities in MS DOS 6.0, *PC Magazine* (April 13, 1993) noted that the inclusion of MSAV in DOS 6.0 is "a serious blow to the street-vendors' antivirus software market." This article, however, necessary all this antivirus software really is, given the extreme non-incidence of verified virus sightings. But certainly this is a very active part of the PC software market.

MWAVE.LIB from Central Point's Anti-Virus for Windows relies on many dynamic link libraries, including MWAVMR.DLL (FAT manager), MWAVSCAN.DLL (virus scanning support), MWAVMSI.DLL (absolute disk I/O), MWAVDIR.DLL (dialogs), and so on. That DOS 6.0 now includes Windows DLLs raises an intriguing question: Are these DLLs now part of MS DOS? The names of the functions these DLLs export are possibly visible using utilities such as Microsoft's EXE2DEF and Borland's EXE2IMP (calling functions in DLLs is discussed in upcoming detail in *Undocumented Windows*). For example, MWAVSCAN.DLL exports many named functions, including CPAY, Scan for Viruses, CPAY Scan root for Viruses, CPAY Kill File Virus, and CPAY Kill Root Virus. Are these now part of DOS? Should this be documented? What the heck knows?

It is good to know, of course, that none of these Windows executables or DLLs seems to use undocumented Windows APIs. Naturally, MWAVSCAN does lots of low-level stuff (making heavy use of DPBI, for example, in its modified SetSelectorBase and SetSelectorLimit functions), and so on. One of the oddities of MWAVMR.DLL, early on, at least, contains calls to INT 2Fh (AX=0282h and AX=54h). A modified INTRPL (these functions are part of the PC tools interface. Given the current existence of Microsoft and Central Point, or at least the incorporation of so many former Central Point utilities into DOS 6.0, should this API now be considered part of MS DOS?

Of course, Windows also uses DLLs. MWAVMSI.DLL, a collection of functions for absolute disk I/O, also exports about a dozen named functions, including AIO GetDDBHeader, AIO GetLastEntry, AIO GetNextEntry, and AIO GetNextList.

One function provided by MWAVMSI.DLL is AIO GetNextList. Not surprisingly, this function calls INT 2Fh (AH=54h, AIO GetNextList) uses the Windows DOS3C API (function to do the INTRPL). For present concerns, functions in this category of any use are needed, but Microsoft needs to know its INT 2Fh (AH=54h) list are, at least. There's no question of insider knowledge. Note that the register names of this list are exactly the same as the RDI (owner's interrupt list) and general PC talk one can find in the MS DOS source code, which takes this data structure, `SystemVary`.

Another example is provided by MWAVMSI (in File 130) also calls INT 2Fh (AH=52h) and creates a protected-mode pointer to the SystemVary SFT (Figure 1-20 shows the code). The code converts the real-mode SFT segment to a protected-mode selector using DPBI function 0002h, which creates a per-processor selector, but with no responsibility of time. While DPBI itself is hardly undocumented, the DPBI specification is available only on exchange from Intel. Microsoft barely documents the presence of DPBI (see "6.9.1.1.1 Windows for Windows API Windows SDK contains a total of 100 pages on both the DOS kernel and DPBI "Windows Applications with MS-DOS Functions" *Programmer's Reference*, *Volume 1*, Chapter 29), plus skips chapter, which also incorrectly claims that Windows supports a system that DPBI lists a set of seven DPBI functions that Microsoft approves for use by Windows applications. Function 0002h is not one of the list. Microsoft's rendition of the golden rule would appear to be "Divas I say, not as I do."

**Figure 1-20: MWAVMSI.DLL Creates a Protected-Mode Pointer to the SFT**

```

7 087A mov ah,52h ; most popular undoc function
7 087C call far ptr 0053CALL ; call INT 2Fh
7 0881 mov ax, es:Cbx+43 ; sysvars[63]=SFT chain offset
7 0885 mov word ptr SFT_PTR, ax ; save offset
7 0888 mov bx, es [bx+6] ; sysvars[6] > SFT chain segment
7 088C mov ax, 2 ; DPBI Segment to Descriptor
7 088F int 31h ; input BX=real mode segment
7 0891 mov word ptr SFT_PTR+2, ax ; output AX=mode selector

```



MWAA ABSTRACT 111 makes extensive use of this protected mode pointer to the `NTL` chain. For example, the exported `MDL UpdateSFTInfo` function of course uses `SFT_POINTER`.

MWAA ABSTRACT 111 raises yet again (though doesn't answer) the questions of whether it is right for MS DOS to use undocumented functionality in products that compete with third parties' (read *PC Magazine's* statement about the above MSW) way going to deal to the anti-virus software market—and whether Microsoft needs to document the APIs exported from DOS incorporated by MS DOS 6.0. In any case MWAA ABSTRACT 111 constitutes further proof that undocumented functions are used and that it is crazy for Microsoft not to document INI 21h AH=52h C calls or nonopolistic.

It'll isn't taken by Microsoft to use undocumented calls within the context that ship as part of MS DOS. In the case of SHAREDIN, SIBIND and so on, the answer is obviously that it is okay. Every operating system must have undocumented functions. If every interface were documented, then the operating system would be doing its job of shielding developers from complexity.

But what about a utility bundled with the operating system, such as Microsoft Art View or Double Space, that compete with already existing third party products. Here the answer is less clear. *PC Magazine* (September 14, 1993) has complained that “in adding these utilities, Microsoft has clouded the definition of exactly what a operating system should do.” Many would disagree with this statement. What is separate from the operating system? The very existence of third party products would seem to be the key test. Consumers have demonstrated a desire to purchase in large numbers utilities that compete with software that comes bundled free as part of MS DOS. Executive Microsoft CD ROM Extensions, MSC DEN, have competitors—a product from Local Software. Where does the operating system stop and the third party utilities' (the market) begin? At what point does Microsoft's use of undocumented DOS calls cease to be operating system internals and turn into an advantage over competitors. At what point is the advantage unfair?

The answers seem to be unclear whenever we get to products such as Windows for Workgroups and other Microsoft products that are sold separately from MS DOS. But even here there isn't a simple answer. If Microsoft bundles Windows with MS DOS, as it would with one of our sporting “Chicago” operating system, does the problem go away? Is bundling a good solution to the problem of using Perlths. On the one hand, the continued existence of products such as Desktop Textbox to the desire that it cost some consumers have to buy pieces of the operating system separately, somewhat like the aftermarket for automotive spare parts.

In the end, it is unclear how clear today Windows and DOS are separate. Windows relies on INI 21h function 52h, SVA, the CDS, the SFT, and other undocumented DOS calls in its structure. Microsoft's failure to document these, or to provide documented equivalents, is monopolization with no benefit to consumers and with harm to third party developers.

## No Problem?

The problems with undocumented DOS are that 1. some of these functions and data structures are absolutely critical to DOS programming, 2. Microsoft uses these functions and data structures in its own software that competes with third party utilities, and 3. programmers have repeatedly had to request that Microsoft document its structure. While Microsoft did finally document a number of previously undocumented DOS functions when they released DOS 6.0, there are still many INI 21h AH=52h in particular, not to which Microsoft denies there is any problem.

Free software Microsoft issues, as changes over time, things. On occasion, Microsoft representatives will deny even that undocumented DOS or Windows or OS 2.0 are available.

But usual, what Microsoft denies is that Microsoft applications use or derive any advantage from these calls. For years, Microsoft has claimed that, between its operating systems and applications groups, there is “separation of church and state” or, using a term from the insider trading scandals of the 1980s, a “Chinese Wall.”

Stephen Marics and Paul Andrews discuss the Chinese Wall issue at length in their fine biography of Bill Gates: *Gates: How Microsoft Made Revolution an Industry—and Made Himself the Richest Man in America* (1993). There have been other biographies of Gates and histories of Microsoft, but this is the best and most thoughtful one. I was especially struck by a *People* magazine story report on Gates' life, who cares—but rather a—depth study on how Microsoft does business. On the question of the supposed Chinese Wall between Microsoft's two businesses, Marics and Andrews report:

As late as 1984, Microsoft executive VP Steve Ballmer had stated that "We have shown it is not just that there is a very clear separation between our operating systems business and our applications software. It is like the separation of church and state. And if you do it plain—straight—you can't expect to get the business." But except for Ballmer's favor to pursue "get the business," it was hard minimalism. There had been no formal division between systems and apps until mid-1984.

Interesting is a lot of time years about Microsoft's applications conflicting with Microsoft's rise as the father of the operating system wars, from that time of OS/2.

In 1985, columnist story about OS/2 was more candid. *InfoWorld* columnist (now *PC Magazine*) columnist Steve Marics reported that several developers "I've talked to recently believe that Microsoft has an unfair advantage in developing applications that will run on OS/2. One has even whispered the ugly word 'antitrust'."

Microsoft's Adrian King stated that "There are no undocumented interfaces in OS/2," thus elevated its position that "If there are any undocumented interfaces to OS/2, they are for operating system's hardware." Finally King admitted what was patently obvious: There were no other rules that truly separated the applications and system sales of the company.

So there too, Microsoft's claims about the claim that there were no undocumented interfaces in the case of OS/2, it was more absurd. A good example is the `DosPTrace` function most of systems programmers had to use to get IBM or Microsoft would document—see "Stalking Cd' Faults," *The Daily Journal*, February 1990. As good as not documented functions never die, it is amusing to see the Microsoft make the link for `DosPTrace` to create the important Windows WinDbg debug function was not only Microsoft's own to document, supposedly because the function is "generous," which doesn't in fact appear to be happening.

Remaining is the supposed separation of Microsoft's applications from its operating systems and the "get the business" of the Chinese Wall, as a particularly revealing and balanced summary:

It also seems that a clear separation would be readily decided and the "ugly word" antitrust, the saying spoken. "There's a kind of a wall but there is not a watertight wall between the two Windows developers said. It seeps when you get to Bill. And a developer said something like there was going to be a little bit friendlier to the guy down some more something that the group in Boston." If a Microsoft app developer walked down the street in Windows developer and asked for a favor to make Excel run better, was that [violation of] antitrust?

Besides in the age of the great worst of Microsoft—just asking for something didn't mean you didn't Windows developer Steve Wood—kept begging the languages group to give its C++ compiler some cut, maybe to be more desperate missing for Windows developers to be honest and out. The languages guys didn't care—and it was years before his wish list was implemented (pp 349-350).

It is also who has ever dealt with the Microsoft languages and operating systems groups, this part rings especially true. There appears to be an intense rivalry, and something approaching intense dislike, between these two groups at Microsoft.

Unfortunately, though, the Chinese Wall didn't hold much water because—as was shown in *Undocumented Windows*—whatever separations did exist was completely arbitrary and obviously not a matter of great importance to Microsoft. Why, in fact, somehow Microsoft applications such as Excel and Word for Windows could (1) send and receive mail, Windows "network" via through Mike Maples, for example—admittedly *InfoWorld* (December 30, 1992) said "the biggest worry would be if we were using secrets in undocumented things and we very consciously avoided it." It is important to stress it as having its applications—its undocumented functions—was not a matter of even a minor part of Microsoft's competitive strategy, but nonetheless the company's intentions to have some deliberate policy of not using these functions was shown to be totally bogus.

Having failed to deny convincingly that undocumented functions exist in that Microsoft uses them, Microsoft's third defense was that some "secret" functions exist and that Microsoft applications use them, but the functions were not for such purpose. Far from giving Microsoft a competitive advantage, so this argument goes, calling these functions actually puts it at a disadvantage. Obviously, I don't see how. We have always been at war with Heaven.

This approach came out during the "radio press tour" created by *Undocumented Windows* (see, for example, "Microsoft Runs Red Over 'How About Windows' With *Street Journal*" (September 1, 1992); Microsoft ends its frantic, its aggressive public relations (i.e., *Windows* Edition) press secretary press releases on the topic "Microsoft Statement on the Subject of Undocumented APIs" (August 31, 1992) "Questions and Answers About Documented and Undocumented APIs" (undated), and "Undocumented Functions" (undated) which put forward this idea that calling undocumented functions puts Microsoft at a disadvantage.

Take, for example, the undocumented Windows function GetTaskQueue, which *Undocumented Windows* shows that Microsoft QuickC for Windows uses (as in the PC marketplace computers are applications). Microsoft's "undocumented functions" statement also describes its status as "Not available in beta releases." I wonder where the Microsoft statement said that "why GetTaskQueue" "is a bit of disadvantage" which is the same thing I say to every other undocumented function that Microsoft was caught calling.

While some of the undocumented functions that Microsoft applications use are *really* useless, representing obsolete code from older versions of Windows, this demand was not the case with GetTaskQueue. Disassembly of QWWINI shows that the primary way this undocumented function is determined if there is a message posted for the user's program, coming within the QWWIN integrated development environment. This is QWWIN needs to determine whether the user's program has executed and the screen has startup code that needs to be trapped to establish a task queue. If an application desires GetTaskQueue, sending it messages containing `WM_SETTASK` or `QMWIN` contains the following code:

```
if (GetTaskQueue(hTask) != 0) // If there's a queue
    PostAppMessage(hTask, ...); // send a message
DirectedYield(hTask); // either way, do DirectedYield
```

Microsoft claims that the function SetMessageQueue is a documented alternative to GetTaskQueue, but whoever wrote that program and looked up the word "queue" in an index SetMessageQueue simply gets by way of a message queue, it is no help with QWWIN's problem which is to ensure that another task has a queue. But someone—my wonder who—who designed this function is such a good character to be undocumented one. QWWIN uses the undocumented one.

There are other complex puzzles GetTaskQueue— but this is supposed to be a book about DOS, not Windows. It's hard to keep them separate these days, so we'll move on. Microsoft's "we call these functions as a nice neighborhood way of putting ourselves at a disadvantage" is their own thing, the playing field is again basically not very effective.

We expect Microsoft's next defense to be that sure, Microsoft uses these functions, but so do they are useful, but that is wrong, uses them. In fact, one of Microsoft's statements is that "it's okay

Michael Williams controversially notes that “MS-DOS has a number of undocumented APIs that are well known and understood and used by ISVs” (independent software vendors). Well okay, but then why not document them?

## DOS Documented

We had as yet to be subject of why Microsoft leaves critical pieces of DOS functionality undocumented in a manual, but it is important to note first that the company has to its credit finally documented some of the things we know and well understood previously undocumented functions.

In the summer of 1991 Microsoft Press issued the *Microsoft MS-DOS Programmer's Reference*, which includes for new INT 21h and INT 2Fh functions in MS-DOS 5.0. This book came out after the publication of the first edition of *Undocumented DOS*, and the Microsoft Press advertising material for the program's reference was “DOS Documented.” Microsoft released a nearly identical version for MS-DOS 6.0 in the summer of 1993.

Unfortunately, the programmer's reference did document some previously undocumented calls. Most of the Microsoft-documented the following previously undocumented INT 21h functions:

- Function 11h (Get Default DIB)
- Function 32h (Get DIB)
- Function 34h (Get InDOS Flag Address)
- Function 4B01h (Load Program)
- Function 50h (Set PSP Address)
- Function 51h (Get PSP Address)
- Function 510Ah (Set Extended Error)

The very documentation for these old functions is unfortunately far from complete. One gets the sense that it is a stretch to claim that this reference for Microsoft was simply to claim that these functions are now documented. The new documentation for the previously undocumented INT 21h functions is the following: precedes:

- Function 11h/32h: The programmer's reference claims that these functions are only for DOS 5.0 and higher. In fact, these functions are present all the way back to DOS 2.0. This is an important piece of information because a large number of DOS 3.x installations are still in use, and most PC disk utilities which rely on function 32h must be able to run on these machines. The DIB structure provided is only accurate for DOS 5.0 and higher. Basically, if you believe Microsoft's documentation on the disk utilities can only be written for DOS 5.0 and higher, which we know not to be the case.
- Function 34h: The documentation does not mention the critical error flag located in the `dwCriticalErrorFlag` DOS flag. As we saw earlier, the Windows DIBMMGR relies on being able to decode the error value from function 34h to get both the InDOS flag and critical error flag into a single word.
- Function 4B01h: The documentation for the TOAD structure incorrectly reverses the order of the `dwIP` and `dwNSP` fields. Some errors are of a nature unavoidable, but the identical error had appeared earlier in the book *Developing Applications Using DOS* by three IBMers, Ken Christophs, Barry Engelenbaum, and Shon Sager, 1990. Presumably, Microsoft used this book as source for the source material for the MS-DOS 5.0 programmer's reference. The MS-DOS 6.0 programmer's reference does correct this error. For a full explanation of function 4B01h, see John Paterson's “The MS-DOS Debugger Interface” in the first edition of *Undocumented DOS*.
- Function 510Ah: This function is documented only for DOS 4.0 and higher. In fact, it is present in 3.x and higher. Furthermore, because the manual does not mention the other AH=50h

functions. The `ERROR` structure has a field that will, in best faith, catch readers. According to the programmer's reference, only I/O "identifies the computer or errors that occur on remote computers. It is not a system call." Without documenting it (see 5.100), Server Function Call identifies DOS programs for structure, so it is a special case. It is a system call, but not so. The call ID is, in fact, a new status IDOS's USER ID field, discussed earlier in this chapter. Under Windows I must not mock, but had thought this corner the VMID, which a DOS program is free to use, by INI file V.1688. It does, as expected, enough to indicate that Microsoft's documentation of this function is correct. The MS-DOS 5.0 programmer's reference also stated that the `ERROR` structure goes to `OSM`, which, in fact, goes to `DS:DI`. This, at least, was based on the MS-DOS 6.0 programmer's reference.

Well, maybe you're a chameleon, and it is good to have these functions finally brought into the official INT 21h fold.

In addition, Microsoft documented INI 283, MS-DOS file Header, and a number of INT 21h calls:

- Function 0100h: Get Control Device (the 1st edition of *Undocumented DOS* incorrectly identified this call as PRINT.COM Check for an Output Device)
- Function 0600h: Get ASSIGN.COM Installed State
- Function 1000h: Get SHARE.EXE Installed State
- Function 1100h: Get Network Installed State
- Function 1400h: Get NLSFUNC.EXE Installed State
- Function 1A00h: Get VSS/VSI Installed State
- Function A300h: Get KEYB.COM Version Number
- Function A380h: Set KEYB.COM Active Code Page
- Function A382h: Set KEYB.COM Country Code
- Function 19000h: Get C.R.A.P. Add-Ons Installed State
- Function B700h: Get APPEND.EXE Installed State
- Function B702h: Get APPEND.EXE Version
- Function A704h: Get APPEND.EXE Structure List Address

Two interesting points emerge from this list: Microsoft didn't state that I/O was an INT 21h function 34, 4001-50, and so on, with, according to the documentation, have existed as far back as DOS 2.0. It is clear, should there have been any doubt, that they really do not depend on undocumented DOS calls, and usually ones at that.

Second, if you're using these undocumented functions, you may not notice anything, or being change for now, or clear sharing opportunities in other situations. They were documented. Of course, we only saw this 100% ratio of our set—respectively, 100%—which was expected, and we mentioned functions, so we do not need to be too wide. However, the last two undocumented functions, `SET` and `SET`, are not so easy. We did not see this, but just a short.

The *MS-DOS Programmer's Reference* lists almost all of the user and kernel-mode DOS functions in their original state. Of course, it is not so easy to find even more so. The core `GET` and `SET` functions, for example, are not in the `OSM` or `OSM` structure, but the `OSM` structure is documented, under the name `ARENA`. But now, we do not know of the `ARENA` variable, which the `SET` and `SET` functions are in the `OSM` structure. As some other, only one INI 21h API. Discussion was commented. The network-related calls, but there were, and the only address-related documented function was INI 21h V.11000, in which INI 21h interface is absent because the single, or description "Network Call Sequence." As it all aside from documenting some of the new functions, DOS is not the programmer's reference, did not add anything to the state of DOS as we know it. The DOS 6.0 version at least includes the `OSM` structure, `OSM` and `MRE` for operation in the `OSM` structure. To look merely put Microsoft's goal, a 2nd belated approval on a similar ready, or `OSM`, just for it. As

Page 8 Book 5, part 1 at the link to the semi-independent *Microsoft Network Journal* January February 1992: "We know that Microsoft seems to have finally realized that a number of 'undocumented' functions were not documented anywhere but the MS-DOS technical reference."

## Why Leave Functionality Undocumented?

Why do let Microsoft do a better job with the DOS programmer's reference? Why, for example, give the programmers obvious information DOS programming has it still not documented IN 21h function 82 (see the list of lists). Why are they still heading back on the network redirector specification? Why do they look for it in IN 2A critical section stuff?

It seems so. Microsoft's reasons are not really very different from the reasons other vendors document their systems: users can't everthing that a person would ever need to know. The difference is simply that Microsoft now goes such a central unit in the software industry that when Microsoft fails to document something important, this failure affects many more programmers much more deeply than say, I suppose, a manufacturer of DME construction materials (1/3).

Of course, it is also a documentation's simply a resource allocation problem. It is not easy finding qualified technical resources, and those who are qualified would probably rather write code. Any company that cannot afford all-task leading good technical writers can sympathize with Microsoft's problem. It is not a typical Microsoft when it does find good people to write excellent documentation, such as the one for the Microsoft Developer Network (MSDN) CD-ROM. If you have the MSDN CD-ROM, you can in particular see the excellent articles by David Long.

Another major documentation problem was the same reason for inadequate documentation: it puts the company in a difficult position. For a new Light Microsoft release matrix have on several occasions considered the company's desire to set of the key software standards. Microsoft was not even then a dominant player. Microsoft's means such as by phone, copy, fax, and television. There is in each business some with this. However, if Microsoft wants to set every standard and control every API, it has to make sure that it has the resources to maintain these standards and properly document them. Yes.

Another work-out problem with Microsoft is simply that its eyes are bigger than its stomach. It works on a wide standard set that it cannot have adequate resources to document. The problem is, of course, not only resources, but lack of personnel. Microsoft's policy of deliberately under staffing projects is well known. It is not a good work-out. Some time has to "give" and documentation is a non-negotiable task. I suspect when the existing documentation software say 85% of the users it is sufficient to do the "work" necessary, except what is needed by the remaining 15%. There are other reasons for under staffing, just as there are in big lives.

Another thing to be mentioned here for the rest of the industry. If Microsoft's eyes are bigger than its stomach, again Microsoft takes on projects for which it can do only an adequate job. In technical resources, Microsoft has taken to providing 88% solutions; this situation opens up many niche opportunities for other companies. Microsoft's success in the memory management market is a good example of this. As a company's sales declined after the introduction of DOS 5.0 would shortly lead to bankruptcy. This is not a case of "big eyes or vendors such as Qualitas and Quarterdeck. Instead, the vendor's memory management was considered free of charge with DOS 5.0; the market for memory management was simply split. Apple's other memory manager's know how to "hook a gift horse in the middle" and Quarterdeck's *MemMan* DOS 5.0 legitimized the deal of 386 memory managers, but Microsoft's *MemMan* did not. It is a classic case of memory manager. Companies such as Qualitas and Quarterdeck were able to export their contributions. This book is itself a good example of how Microsoft's tendency to provide only partial solutions can create opportunities for others.

Another reason is that some of the things like the network redirector that Microsoft ought to fix are out of the case, and documented as more valuable to it as their undocumented "magic" state. Documentation is a state documentation, but an important undocumented inertia like the redirector is

"technology." Microsoft has made the redirector source code available to selected vendors, apparently in exchange for things it wanted. It is hard to have such technology exchanges when the "free money" is nothing more than a chapter in a book.

There is a detectable pattern—the things that Microsoft doesn't document, as well as to the things it takes many years to document. This includes, as present in DOS 2.0, the three alpha documents (into DOS 3.0 are, for more examples). Microsoft consistently underestimates the need that programmers have for low-level information about its operating systems. The "anonymous" Microsoft attitude seems to be that if you aren't working at Microsoft, you shouldn't be messing around with SEI chains and CD pointers to the first place. What does it want to know that for? Of course, many people outside Redmond don't share this attitude, and even at Microsoft, the "break-down-and-document" at least some of the low-level interfaces that a lot of developers need. Microsoft may be showing some new openness in its Win32 SDK, which includes the source code for some key utilities, such as DIALOG, PYDWR, SPY, and WALKER. Source code. What a concept!

Another reason Microsoft doesn't document some important functions is that, as one developer at Microsoft once put it, everyone at the company is "too damn busy" doing it. Sure. With third-party vendors need something that works on machines running DOS 3.0 today, developers at Microsoft are already being "out and preoccupied" with DOS 7.0. (Oh, sure, it's also true that's going away in DOS 7.0!)

Microsoft's attitude seems to show it wants to document things that it knows or thinks it expects or plans to go away in some future version of the operating system. It is common also that the company is so far ahead of the game that it does little good to the third-party vendor who may not have perhaps a million uses, yes, but a common and need to accomplish some very real work today.

From Microsoft's perspective, looking at the low-level functionality of many programs, it seems perfectly sane to reserve the names of DOS and to tell developers that if they would honor that, then both these areas and use the other programs might or might not work reliably, clearly. Microsoft has issued a policy statement about programs that use undocumented DOS functions and other resources. Reprinted in the Use of Undocumented MS-DOS Entries (Q14 01, September 5, 1988).

Microsoft does not give out the exact source code, and documents a system, like it likes it. It calls things on contracts to be understood, but it says, they are not supported, so you give NO guarantee that they will exist in future releases of DOS. If you're serious about these functions, having a discussion with us, and organizing them into a company contract, there is a real potential that your application will not work in future DOS releases. We strongly advise against using undocumented features for these reasons, and we have no information about their use.

Once a software developer documents some feature of a product, the developer is more obligated to support that feature in future releases. Microsoft has several problems regarding contracts of MS-DOS that it has already commented. The existence of some CP/M compatible machines, such as the Compaq Disk II, the and the structure of the DOS Program Segment Library (PSL) are good examples. The ICB is a good example of why operating system standards should not be exposed to applications. Documenting, but they can make a range of them difficult and can thus create many answers.

So far, Microsoft's actions pertaining to document key interfaces and low-level information. Now, it is seems particularly vicious, conspiratorial, or even the product of malicious, thoughtless, wanton behavior.

But, its passive-aggressive practice of neglecting to document key interfaces that it knows it does have broader implications. The problem isn't "conspiracy," which seems to be the word that many analysts used to describe the notion that anything could possibly be so simple. Redmond's idea of a conspiracy sounds a bit far-fetched. But the problem isn't "conspiracy," it's simple. Given Microsoft's 95% market share in PC operating systems, this sounds downright impossible. The problem is, of course, we all know it. Why should Microsoft be any different?

## Documentation and Monopoly

What does Microsoft's position as monopoly supplier of DOS have to do with its failure to document key pieces of DOS? As one I earlier, the computer trade press may have wrongly focused on possible benefits that Microsoft applications derive from this arrangement and failed to note the benefits that DOS itself derives from undocumented interfaces.

It's hard to say how far Microsoft managed to keep such a restrictive market as DOS almost entirely to itself. Simple economics would tell us that given the relative simplicity of the DOS interface and given the huge volume of money to be made by taking even a small piece of the DOS market away from Microsoft (see Chapter 4), Microsoft would have more competition for DOS. Why aren't there at least three or four other companies producing high volume DOS workalikes, in the same way that many companies produce word processors. Why is the competition for DOS so thin?

One explanation is that the operating system might constitute a so-called "natural monopoly," that is, a market in which adding competitors drives prices up rather than down. Microsoft itself may believe that DOS is a natural monopoly. Manes and Andrews' biography quotes Bill Gates at a Rosen Research Personal Computer Forum back in May 1981, discussing how volume and standards might lead to a natural monopoly:

We do set de facto standards. It's only through volume that you can offer reasonable software at a low price. Standards increase the basic machine you can sell into a market. I should not say this, but in some ways it reads in a video data product except for a natural monopoly, where somebody proposes documents, property trains, property, makes a particular package and through momentum, user loyalty, reputation, sales force, and price builds a very strong position within that product.

Gates' point is understood: the connection between volume and quality earlier than anyone else in the PC software industry, which industry of course he largely defined in the first place. Still, it is somewhat surprising to see Gates speaking of monopolies even back in 1981. If the PC operating system is a natural monopoly, shouldn't it come under some form of government regulation, just like telephone service, electric and electric utilities. The idea that Microsoft might consider the PC operating system a natural monopoly is particularly worrisome given the company's clear trend of putting more and more functionality into the operating system, particularly Windows, that was preciously thought of as proprietary. If, after all, the Windows GUI 2.0 specification is a good example, does this mean some protection on established boundaries?

But even if we see how there is an "artificial" monopoly here, there are certainly many different ways to implement the fundamental DOS interface we get to the undocumented parts of the interface, and so on. Still, the entrance of some strong competitors into this arena would drive prices down, not up. While developers would surely rightly not welcome the presence of multiple DOSs or multiple implementation variants of Windows, it would be a testing time, not, at least at first, overall competitors would benefit. The highly competitive market for DOS—just as competition from AMD and Cyrix and now perhaps Motorola's PowerPC—has helped Intel speed up its product development cycle, so it would competition for DOS make Microsoft more innovative.

DOS's lack of new software, not to mention largely because of the lack of competition. And what little software there has been since MS-DOS 4.0 has come from changes made in DR-DOS (see Chapter 4). Of course, some of this relative stagnation in DOS is based on the need for backward compatibility. But if there were more competition for DOS in the same way that there is for word processors and spreadsheets, Microsoft's "Chicago" would probably have been out by now. With little competition, Microsoft was able to stave off DOS and instead devote its efforts to exercises such as NT. The less competition it had for DOS, the more Microsoft was able to ignore DOS and simply milk it as a "cash cow."



The Microsoft near-monopoly on DOS provides little benefit to consumers, so it's not a natural monopoly. But then what explains the lack of competition in this "creative market" — the incentive to developers if there were more than one important DOS is one explanation. Microsoft's per-machine rather than per-copy OEM pricing of DOS is another. But an additional explanation — and one which will return us to the subject of undocumented interfaces is the fear of incompatibility. A key to Microsoft's near-monopoly over DOS is the notion that the only power, DOS is MS-DOS. Because DOS itself is another small piece of code which has been highly leveraged into an incredible business — and because implementing the documented portions of the DOS interface would be relatively simple, there are probably infinite ways to implement IN 1.21 + MS-400 without infringing on Microsoft's copyright. Microsoft's DOS monopoly depends largely on this sort of "compatibility."

But compatible with what. It is relatively simple to be compatible with the documented DOS interface. Failure to do so spells a big "No" what does this "compatibility" mean?

In part of course, compatibility is just marketing product differentiation made visible. Microsoft wants to promote the idea that MS-DOS is the only conceivable DOS. In this sense "compatibility" is just another way of saying "functioning like last year's MS-DOS."

But in operating system terms, work as part of an integrated package together with other products. So there is something going on in the issue of compatibility. It really means compatibility with *undocumented* interfaces. After all, being to be compatible with documented interfaces is just a big Microsoft can make — good argument that only MS-DOS implements a true undocumented DOS interfaces correctly — and that therefore, only MS-DOS is 100% DOS compatible. This appears to be the core of Microsoft's defense of the VARD code in Windows, i.e., that 100% DOS compatibility "100% DOS compatible" should not pass the VARD test. As we saw earlier, the VARD test for 100% DOS compatibility depends entirely on undocumented interfaces. Producing a compatible DOS then, is really a matter of getting the undocumented interfaces right.

Because it is relatively easy for competitors to be compatible with a documented interface, companies try to create what Ray Venkes calls "artificial kingdoms" — i.e., selectively documenting only parts of their products' interfaces. Consequently, Microsoft uses undocumented interfaces to shore up its dominance of the DOS market. Whenever an application developer advances a new workbench or adds an undocumented DOS services and uses the same residential DOS, normally successful applications now, by if reverse engineering to use MS-DOS — i.e., to that particular sequence of bytes rather than to the DOS standard interface. Conversely, relying on undocumented calls and structures makes it more difficult for a developer of DOS workbenches.

But for Microsoft to share an — i.e., a thought from undocumented interface — important applications have to use them. If so, could it turn out that Microsoft would in some way actually have to encourage key vendors — especially the ones that do most of the work — Microsoft indeed does claim to help ISVs use undocumented calls when necessary. This has always raised the question of why Microsoft doesn't just document the calls. The possibility of using undocumented interfaces essentially as a means of product differentiation provides one answer to that question.

Indeed, many so-called "undocumented" interfaces are actually "selectively documented" interfaces in which Microsoft has selectively allowed some of its competitors and/or customers access to an interface, while denying similar access to other companies and to the rest of the developer community. There are numerous instances of this, including the network redirector, the Global FMM Import specification, and the IOCTL code for Windows 3.1.

Earlier, it is noted that Charles Wells' already-mentioned "open" claim what software engineering teams sometimes call "standard" or "legacy" that is narrow, well-documented interfaces. But the software engineering tends fail to mention that to properly document interfaces is also to throw them open to potential competition. Conversely, undocumented interfaces are a way of creating and reinforcing a *monopoly standard*.

Choosing a critical standard depends on publishing this standard and having others write to it. Otherwise, it is hardly worthwhile. But to publish a standard opens it up to potential competition from clone workers, re-creators, and so on. Is there any trick that would allow one to own a critical standard, at present or in the future, or at least delay competition? MS-DOS shows that the answer is yes, and undocumented DOS interfaces, on which key applications rely, is part of how the trick works.

## Fear of Undocumented DOS

Why, of course, that MS-DOS includes a lot of undocumented functionality. Microsoft doesn't document these features because it wants the freedom to change, or discard them in future versions of MS-DOS, a choice seen in the cases just discussed. These features may be more useful to Microsoft in their future versions than you are. Armed with *Undocumented DOS*, you may know a lot about these functions and their uses. It is interesting to know that MS-DOS set this file address as an internal variable table if you make INI file function calls. But a session in this stuff is a program.

People who are afraid of talking to Microsoft generally say no, you can't. So do other programmers as well. When you start areas of competition, the use of reserved, undocumented, or unspecified features is a good way to create incompatible software. Use of undocumented features is not generally a part of a specific software engineer's ethical code. It is hard to believe that using undocumented features is one of the ways to write stable and correct utilities for MS-DOS.

There is a certain step-by-step writing style, undocumented DOS, and some programmers have found it easy to take the dogmatic view that programmers should never ever use undocumented DOS functions. It might be a fair way to distribute to others. For example, the author of a well-written web magazine and computer introduction to IBM programming, Thomas A. Wadlow, *Memory Research: The Development of the IBM PC*, 1987, writes:

None of the programs in this book use the INTDOS call, and for good reason. INTDOS is "undocumented" - something that has no meaning. The first is, of course, that you cannot look it up in the DOS manual. The second is that Microsoft, the vendors of DOS, reserve a right to occasionally delete this function from subsequent versions of DOS. In fact, the INTDOS call is shown here, symbolically, under DOS version 2, which was any of the major commercial versions. In DOS version 3, the call still exists, but has changed quite a bit from its earlier version. In DOS 4.0, this function does something quite different, thus calls for the INTDOS function will fail miserably.

For that reason, use of the INTDOS function call or any undocumented DOS function is not recommended.

The DOS reference in the back of Wadlow's book has entire pages with only a note at the top such as:

AH = 05AH (52)

Unsupported  
Universal function

The rest of the page is left blank!

And more than that, a case I found (INI file function 34). I don't have all the information about the page, but in any case, DOS version 3, the next, then calls to the INTDOS function with no other information. But Microsoft, IBM, from the first edition of *Undocumented DOS* used this then undocumented function address, presumably, correctly as DOS 2.x, 3.x, 4.x, 5.x, and 6.x in the DOS manual. DOS 2, the original release, is IBM DOS.

What, what happened with function 34, was that, far from failing miserably, Microsoft discovered that the most useful, widely reserved calls are destined to become documented, then using such functions is not something more dangerous than gas, you a year or two jump on your more expensive computers. As other good examples, yes, the SetCursorSelector() and Base functions in Windows 3.0 were more successful than Microsoft documented them in Windows 3.1. Was it a mistake to see the stack when they were undocumented. In retrospect, clearly it wasn't. There is something of a

Heisenberg effect at work here. The reason that previously undocumented functions became documented is that programmers start using them. The documentation is to an extent market driven.

An informal poll seems to indicate that developers of commercial PC software—software that must maintain a high degree of security and compatibility, sometimes on millions of different machines—are in general less fearful of undocumented DOS than programmers whose work needs to run on only one or two machines. A clear outlier paradox.

One explanation is simple, but it is mostly mass-marketed shrink-wrapped utilities rather than say, in-house databases or critical market applications that require a undocumented DOS. Another explanation is that programmers who work for large commercial software companies can better afford the possible higher cost of working with undocumented DOS. Software that uses these functions perhaps requires more testing and time in the marketplace than that which uses normal DOS code.

In any case, attitudes toward undocumented DOS resemble current opinions about the 'gotta' construct in programming languages. Making a professional programmer recognize that 'gotta' possibly disguised as 'to go top' is sometimes necessary. Like 'gotta,' programmers should avoid undocumented DOS, but not worry excessively. Software's misbehavior involves tradeoffs and compromises, not fixed logical software construction aspects to be cringing (not religion).

Even though a lot of undocumented API calls is particularly prevalent in widely distributed software packages for the PC, the developers that using these functions is somehow allowed only to the home hobbyist level. For example, *File Norton Windows 3.01: A Post-mortem Autopsy* by Paul Yao and Peter Norton, says

It has been our experience that "undocumented goodies" are interesting to look at but dangerous to include in software that is intended for general distribution.

While this is an excellent book, this particular statement is quite odd. Windows 3.01 as we see it reaches heavily on undocumented DOS calls, and Windows was certainly intended for general distribution. More to the point, Peter Norton's own products, the Norton Utilities and the Norton Desktop for Windows, clearly rely on both undocumented DOS and Windows calls. Well, maybe Peter didn't read what Paul wrote.

## Ain't Misbehavin'

As we've seen, a lot of systems software and utilities for the PC use undocumented DOS. With some exceptions like the Windows DOSMBR.VXD, these programs tend to make use of one or two undocumented DOS calls. This is somewhat like using one's right leg, however. It takes only one undocumented DOS call to change the nature of a program.

Let's say that a system is one or two undocumented DOS calls or one program. What type of program is it or how is it? This chapter on Compatibility and Portability in *Windows Advanced MS-DOS 3.0: An annotated registry* categorizes MS-DOS applications by degrees of compatibility. Dumb, in that it provides a checklist, is that use undocumented DOS to the innermost circle of this DOS inferno.

It should be apparent that all rely on undocumented MS-DOS function calls or data structures: a reception of MS-DOS or ROM-BIOS interrupts or direct access to mass-storage devices (bypassing the MS-DOS file system). These programs tend to be extremely sensitive to their execution environment and typically must be "adjusted" in order to work with each new MS-DOS version or PC model. A really old popular font and may be old, but ENR (i.e., EX) network programs and disk repair optimization packages are in this category.

The most important sentence here is the last one. If you write "I believed" DOS applications, you are a good company. Indeed, the purpose of *Undocumented DOS* is to show how you too can write all whatever programs such as Microsoft Windows, the Norton Utilities Stacker, and QEMM. All these programs also tend to be extremely sensitive to their environment. Some of them do have to be "adjusted" to work with a given MS-DOS version. Start using undocumented DOS, and that will be true of your software as well.

These problems already are a fact of life in the MS-DOS world. In fact, using undocumented DOS is one of the surest methods and habits as the standard practice of bypassing DOS and writing directly to the hardware. Windows was supposed to eliminate that problem, but as *Undocumented DOS* has shown, it traded that problem for another one, more sophisticated interfaces.

There are many undocumented functions and data structures for many important tasks. It is your responsibility about MS-DOS than it does about any sort of standard recommended engineering practice. Before we start knocking MS-DOS, though, let's not forget that, if for no other reason than for at least a few on the grounds of the DOS's success, DOS has succeeded in a way that no other, supposed better, operating systems can match. DOS, with all its warts, is an inescapable reality. Using undocumented DOS can not hold a place in any software engineering curriculum, but it is a good exercise in applying engineering core principles to the real world.

Having said all that, of course, what we are a salvage of good engineering practice as we make our way out of undocumented DOS. This book presents many techniques for using undocumented DOS in a responsible, sane, and reliable manner. Some of the techniques recommended in this book are:

- Regularly checking of the MS-DOS version number—not so easy as DOS 5.0 and higher (see Chapter 2).
- Verifying the basic integrity of undocumented DOS estimates by performing an undocumented DOS call and comparing its output with a known value.
- Computing structures (e.g., dynamic data table, check for sizes computed from the DOS version number).
- Never use an undocumented function or structure when there's a documented alternative. Don't be so sure about it. Looking for the documented alternative, don't just assume there isn't one.

It is your responsibility about undocumented DOS as obligated to do a couple jobs of DOS version checking, error checking, and basic sanity. Looking that many other programs that otherwise play by the book. One can not be too sure about it. This is the simplest thing I've ever heard. All programs, whether they use or undocumented features or not, are to do rigorous error checking. Most of the programs in *Undocumented DOS* and its accompanying disk and work properly on MS-DOS versions 2.x, 3.x, 4.x, 5.0, and 6.0. Some of the programs have supported or protected mode using DOS extenders. Many have been tested under different configurations including Windows, SHARE, Dosqview, QEMM, and DnMux with various DOS components loaded in a configuration.

Many of the programs in this book have also been tested under environments such as the DOS comparison, some found in DOS 3.1x and 3.21 and Novell's DR-DOS, see Chapters 3 and 4. These customizations may be too much to ask for, but it is important to gauge the quality of the support for undocumented DOS because any support they do provide a *computer environment*. Unlike versions of MS-DOS itself, which may support one or another undocumented DOS feature simply out of inertia, these supported DOS environments can only support a documented DOS function call or data structure if someone consciously puts it there.

So, there is a large collection of popular PC software that uses undocumented DOS. Are the vendors of all these programs going to get burned with the next version of DOS? It is instructive to read what Microsoft's "OS, Inc." Chief Architect for System Software said about this issue: Gordon Letwin, *Inside DOS*, 2/1988.

It may seem that if a popular application "pokes" the operating system and otherwise engages in unsavory practices that the authors or users of the application will suffer because a future release (such as OS 2) may not run the application correctly. To the contrary, the market dynamics state that the application has now set a standard and it's the operating system developers who suffer because they must support that standard. Usually, that "standard" operating system interface is not even known; a great deal of experimentation is necessary to discover exactly which undocumented side effects, system internals, and timing relationships the application is dependent on.

In other words, when popular applications use undocumented DOS, ultimately it is Microsoft which suffers the inconvenience, not the application's developer. As noted earlier, Microsoft may also derive some benefit in that the popular application has locked itself into MS DOS. The real suffering may be done by vendors of would-be DOS workalikes. In any case, smaller developers, mean while, can ride the coattails of the larger developers. For better or for worse, if enough important applications use undocumented DOS, yesterday's undocumented hack becomes tomorrow's standard. The market has spoken. Amen.



## Programming for Documented and Undocumented DOS: A Comparison

by Andrew Schulman

Let's pretend we work in the installation software group of a consumer software company. For some reason, we're being asked to produce a setup binary that will come from a DOS batch file, retains the number of logical drives on the system, also number corresponds to the `F`ASTDRIVE statement in a recent `CONFIG.SYS`. Perhaps the company is installing software on a Novell NetWare 3.x environment, since `F`ASTDRIVE determines the starting letter for network drives.

The batch will be called `F`ASTDRIVE.C and the ideas that were the batch carry over to DOS, I should return a number corresponding to `F`ASTDRIVE. For example, if `F`ASTDRIVE = 1, then `F`ASTDRIVE.C should return the number five, by which time other DOS utilities that refer to the number six may assume more than five partitions, but since it can be detected (using the `FILE_ATTRIBUTES` MS-DOS command) that's harmless.

The `F`ASTDRIVE utility should also store a file system's `F`ASTDRIVE = 1, but in such a way that reducing the amount of space on a NTFS volume does not cause any output. For example, to make sure the user can see a logical drive, the `F`ASTDRIVE.C program should work in the face of a system that contains other programs, some of which may cooperate work that will usually be our concern, but may incorporate into the following line:

```
echo off
set need6=bat
lastdrv = nul
if errorlevel 6 goto end
echo Requires at least six drives
end
```

Now how do we write `F`ASTDRIVE.C? From C, it's OK, use `CONFIG.SYS` and `ENV` macros. The `F`ASTDRIVE statement's error handling varies from the fact that `F`ASTDRIVE didn't make its appearance in the DOS 8.0, and that its use is optional because 1 is the default `F`ASTDRIVE. We could have a program that once we write a `CONFIG.SYS` file, could verify the way with which the system is booted. It also appears to be impossible to locate an boot drive relative to MS-DOS partition version 4.0 (even versions of DOS 5.1.2) in function 33 (5) since it's this value.

If we're setting up a high-level programming language like C or Pascal, it's a trick that the company's submission binary systems with a function that returns the number of drives. Eric MacNeil's C++ is the best one, from `gpc++` and `gcc++` as well as `FUNCTION_P` escape both of which return a number of logical drives. I don't know about Pascal disks.

More, it's that we were asked to return the value of DOS `LASTDRIVE` (which is the number of the PC's ROM BIOS does not meet the functional specifications for management functions for the `bootsect` file). Logically, we should include RAM disks, network drives, CD-ROM drives, tape back

ip ones and the like. Logical, in other words, means both physical drives and *logical drives*. As shown in Chapter 8 on the DOS file system, much of DOS's extensibility comes from the ability to have drive letters assigned to things that aren't really drives at all. In other words, a logical drive is an MS-DOS construct having little to do with PC hardware or the ROM BIOS. How are we going to write the `LASTDRV` program?

Using the example of `LASTDRIVE`, this chapter looks at how to incorporate the information in the rest of the book into working code in C, 80x86 assembly language, Turbo Pascal, and BASIC. It also discusses in a somewhat general way how to use undocumented features, while showing that certain PC programming tasks absolutely require them.

In some previous chapters, it will also become clear that exploring undocumented features of MS-DOS inevitably requires only a few lines of code. On the other hand, programs that use undocumented DOS features must be more aware of the MS-DOS version number than code that uses only documented DOS. In particular, while undocumented MS-DOS *function calls* have remained remarkably stable through successive versions of DOS to a number (the equally important DOS internal data structures vary) which with each new release of the operating system. Programs that use undocumented DOS must concern themselves with problems in DOS 5.0 and higher; these issues further complicated by the fact that the supposed DOS version number can be modified on a per-application basis. Cross.

This chapter assumes that you want to access undocumented DOS from a DOS program. While this is not particularly unusual, it is not always done, and more developers are writing for protected mode environments like Microsoft Windows. Other DOS programs must still talk to DOS, and sometimes they must access undocumented DOS functions and access undocumented DOS data structures. The next chapter discusses the problem of accessing undocumented DOS from protected mode Windows.

## Using Documented DOS Functions

Before examining how to use undocumented DOS in programs, let's review how to use documented DOS functions. The chapter on documented DOS (we might even say over-documented DOS because some of it has been written about) will pay off when we write programs using undocumented DOS. For now, let's focus on about calling DOS from a conventional programming language; skip to the section on using undocumented DOS. Don't skip this entire chapter though, even if you read a lot of the first section. As you see, we added some code (including a program that modifies over 400 entries) some DOS internal structures, a peek at the disassembled source code for DOS, and a source for using DOS and always a lot of new business involving the DOS version number.

In our sample program, we shall be using `LASTDRIVE`, the first thing to do is browse through a check a look at the DOS program's interface, looking for an `INT 21h` function that returns the number of logical drives.

If you consult Microsoft's official *MS-DOS Programming Reference*, we find that `INT 21h` function `02h` will always return the current disk drive in the system, also somewhat logically, since the two bytes `02h` together with other returns the total number of drives.

```
Int 21h function 02h
Select Disk
Selects the drive specified in DL (if valid) as the default drive
Call with
    AH = 02h
    DL = drive code (0=A, 1=B, etc.)
Returns
    AL = number of logical drives in system
```

In a 3.3 drive `11M2US` and DOS 1.x and 2.x present in more customer sites than you would like to think. DOS returns `AL=2` because DOS supports two logical drives, A and B. These drives hang off the same single physical floppy drive. In DOS 3.x and higher, `AL` returns a the drive code corre-





```

add     al, ('A' - ' ') ; convert to drive letter
mov     dxletter, al   ; insert into string
mov     dx, offset msg ; string in 0510x
mov     ah, 9          ; Display String function
int     21h           ; call MS-DOS
mov     ah, 4Ch        ; Return to DOS
mov     al, 01        ; LASTDRIVE is exit code
int     21h           ; call MS-DOS
main    endp
__TEXT ends
end main

```

You can assemble LASTDRV with any number of assemblers and then link it with any MS-DOS compatible linker using the Microsoft Macro Assembler (MASM):

```

masm lastdrv.asm;
link lastdrv.obj;

```

Or, using Borland Turbo Assembler (TASM):

```

tasm lastdrv
tlink lastdrv

```

When run, the program produces the desired output, for example:

```

C:\UNDOC2\CHAP 2 > lastdrv
LASTDRIVE #

```

## DOS Calls from C

There is a problem making DOS calls using the C programming language. It's not that it is difficult to access MS-DOS services from C; the problem is there are too many different ways to do so. Not satisfied with one technique where a developer changes a word doc C compiler manufacturer for the PC, such as Microsoft, Borland, Watcom, and MetWare, one wonders how much longer the PC marketplace can support so many different good C compilers. Other a wide variety of techniques for calling MS-DOS and ROM-BIOS services. Finding so many different ways to get the same operation is confusing.

If a programmer can't deal with the compilers, however, I think we have to ask who MS-DOS itself has to come out a successful standard means, it is the way Windows does. The only thing approaching a standard programming interface to MS-DOS is the set of functions such as DosGetDrive, DosAlloc, and DosInit that appear in the Microsoft C++ DOS.H header file. Other C compiler vendors such as Borland and Watcom have adapted the DOS.H

Finally, other C++ compilers lack of a standard programming interface to MS-DOS has done nothing to stop MS-DOS's spectacular success. It may even have added the success slightly because it gives programmers one more thing to tinker around with. In any case, we need to discuss a few of the techniques that can be used to make MS-DOS calls from C, including the out86 and in86 functions, more assembly language, and register pseudo variables.

**Int86()** For a long time, the most popular way of calling system services from C on the PC was to use the int86 function which is an Intel 80x86 software interrupt. C compilers for the PC, such as Microsoft C++ and Borland C++, come with a DOS.H include file with prototypes for int86x, int86y, out86, and in86. These functions work with two structures, union REGS and struct SREGS which contain information about the actual CPU registers. To see how functions such as int86x are implemented, see the source code for your compiler's runtime library, for example, [BORLAND]C\CRT\C\LIB\INT86.C\AS:

Listing 2.2 shows LASTDRV.C, which uses `int86()`. This can be compiled with any Microsoft-compatible C compiler for the IBM PC, using either the full version or the command-line version of the compiler.

### Listing 2.2: LASTDRV.C

```

/*
LASTDRV.C -- uses on.y documented DOS, illustrates int86()
Microsoft C/C++: cl lastdrv.c
Borland C++:   bcc lastdrv.c
*/

#include <stdio.h>
#include <dos.h>

main(void)
{
    union REGS r;
    unsigned lastdrv,
           r.h.ah = 0x19,           /* Get Current Disk */
           int86(0x21, &r, &r),    /* call MS-DOS */
           r.h.dl = r.h.ah;        /* r.h.ah now holds current drive */
           r.h.ah = 0x0E,          /* Select Disk */
           int86(0x21, &r, &r);    /* call MS-DOS */
           lastdrv = r.h.dl;        /* r.h.dl now holds number of drives */
           fputs("ASTDRIVE==", stdout), /* output string */
           putchar('A' + lastdrv), /* output drive letter */
           putchar('\n');          /* output newline */
           return lastdrv,         /* return drive number to MS-DOS */
}

```

The C source code in LASTDRV.C is a good illustration of the correspondence between language code in LASTDRV.ASM. One difference is the size of the character string from `sysinfo`: 600 bytes in assembly language and only 5000 bytes in C. Why does LASTDRV.C output the LASTDRIVE string using `putchar('A' + lastdrv)` rather than `putchar('A' + lastdrv)` because while `putchar` can handle 256 individual characters based on byte values 0-255, `fputs` and `fputc` return the number of bytes, not the number of characters, because DOS drive letters are zero-based. LASTDRV.C must subtract 1 from the number of drives to produce the number of the `A` drive, which can then be added to `A` to produce the last drive letter. LASTDRIVE.

**Inline Assembler** A better way to write PC system-level software in C is to use an inline assembler that is put Intel assembly language code directly inside C code. True inline assembler coders often rely on portable C compilers that use `__asm__`. You can't expect MS-DOS or ROM-BIOS calls to work on non-Intel architectures, so this is a perfect place to use some assembly language.

Microsoft C++ and Borland C++ both include an inline assembler. There are a few differences between the Microsoft and Borland versions. A key difference is that Borland does not show `__asm__` in an `asm.h` block. Both Borland and Microsoft put a second-stage assembler right into the C compiler, but Borland uses a second-stage compiler for the two. If you're passing some assembler through to a separate assembler, such as TASM or MASM, allowing you to embed assembly language directives such as `DB`, `equ`, `long`, `org`, `macro`, or `386` instructions directly in your C code. This task is far more difficult with Microsoft's inline assembler.

In LASTDRV2.C (Listing 2.3), note how the preprocessor directives ensure that the compiler can support an inline assembler. Most C compilers identify themselves with a preprocessor definition. Borland C++ provides `__BORLANDC__` and for backward compatibility with the older Turbo C and Turbo C++ compilers, `_TURBOC_`. Microsoft C++ provides `_MSC_VER`, which contains the compiler version number, such as "00" decimal for "0.

**Listing 2-3. LASTDRV2.C**

```

/* LASTDRV2.C -- uses only documented DOS, illustrates inline assembler */
#include <stdlib.h>
#include <stdio.h>

main()
{
    unsigned lastdrv;

#ifdef __TURBOC__
    asm mov ah, 19h           /* C-style comments only */
    asm int 21h
    asm mov dl, al
    asm mov ah, 0x0e        /* C-style hex */
                           /* assembly-style hex */
    asm int 21h
    asm xor ah, ah
    asm mov lastdrv, ax     /* refer to C variables */
#else /* defined MSC_VER || (_MSC_VER >= 600) || defined(_QC) */
    _asm {
        mov ah, 19h        , can include assembly-style comments
        int 21h           /* and C style as well */
        mov dl, al        // and this style as well
        mov ah, 0x0E     , can include C-style hex numbers
                           , or assembly-style hex numbers
        xor ah, ah
        mov lastdrv, ax   , can refer to C variables in _asm
        // Borland would not allow a goto/jmp label here!
    }
#endif

    fprintf(stderr, "lastdrv = %d\n", lastdrv);
    return lastdrv;
}

```

The `asm` means using the `asm` block show the mixtures of C and assembly language that you can produce. You do have to be careful when using inline assembly language. In particular, you must know some important rules about preserving registers. For Microsoft and Borland, the rules are simple:

- You are free to change AX, BX, CX, DX, and IX.
- Inside a function, you can change BP.
- You must always put back any changes to the DI, SI, DS, SS, and SP registers, however.
- Because the compiler has no idea what you're doing with the registers, using the inline assembler turns off global optimizations. In this case, Microsoft C/C++ issues an "inline assembler precludes global optimizations" message.

**Register Pseudo Variables**

Borland provides yet another way to write low-level code—register pseudo-variables. Not to be confused with C register variables, register pseudo-variables map onto the CPU registers but look like C variables. For example, assigning to `AX` is the same as doing a `Mov` to the `AX` register. Listing 2-4 shows `LASTDRV3.C`, which uses register pseudo-variables to implement yet another `LASTDRV` utility.

**Listing 2-4. LASTDRV3.C**

```

/* LASTDRV3.C -- uses only documented DOS,
illustrates register pseudo-variables */
#ifdef __TURBOC__
#error This program requires Borland C++ or Turbo C

```

```

#endif
#include <stdlib.h>
#include <stdio.h>
#include <dos.h>
main()
{
    unsigned lastdrv,
    _AH = 0x19,
    geninterrupt(0x21),
    _DL = _AL,
    _AH = 0x0E;
    geninterrupt(0x21);
    lastdrv = _AL;
    fputs("LASTDRIVE=", stdout),
    putchar('A' + 1 + lastdrv),
    putchar('\n');
    return lastdrv;
}

```

Note that `geninterrupt(0x21)` is not a function call but a compiler directive to emit an `INT 21h` directly into the compiled code. And, although the registers (pseudo-variables such as `AL`) are extremely handy, the code generated by the compiler of course uses the same CPU registers as your code, so you can't rely on values staying in the registers for very long.

### DOS Library Functions

Actually, neither `int86.h` nor `asm` blocks are necessary here. As mentioned earlier, most C compilers for the PC provide a set of functions that map directly onto the most popular DOS functions. Microsoft C provides functions with names such as `dos_getdrive` and `dos_setdrive`. So, naturally, Borland C++ supports these as well as other Borland-specific functions such as `getdisk` and `setdisk` = `FILE_DEVICE_DISK` function set to `0x00000000` with the important exception of Novell NetWare which we discuss later, equivalent to `setdisk_getdisk`.

```
printf("LASTDRIVE=%c\n", 'A' + 1 + setdisk(getdisk()));
```

If functions such as `setdisk`, `getdisk`, or `dos_setdrive` and `dos_getdrive` didn't exist, it would be easy to create them. This is a major use for inline `asm` language.

```

void _dos_setdrive(unsigned drive, unsigned *p_lastdrive)
{
    _asm mov ah, 0Eh
    _asm mov dl, byte ptr drive
    _asm int 21h
    _asm mov bx, p_lastdrive
    _asm xor ah, ah           /* Take 21/0E return value in AL, 0-extend */
    _asm mov word ptr [bx], ax /* into AX, move into caller's pointer */
}

```

```

void _dos_getdrive(unsigned *p_currdrive)
{
    _asm mov ah, 19h
    _asm int 21h
    _asm mov bx, p_currdrive
    _asm xor ah, ah
    _asm mov word ptr [bx], ax /* assuming near pointer */
}

```

### DOS Calls from Turbo Pascal

What about calling MS-DOS functions from other high-level languages? In some ways, it is most simpler to make these calls from other languages, such as Turbo Pascal, because you don't have to worry about which method to use. As noted earlier, having the wide variety of techniques available

not. Ironically, it's so terrible because programmers and writers end up spending too much time deciding which technique to use.

For example, DOS functions from Turbo Pascal requires the DOS unit, which includes the Registers record (similar to a DOS REGs in C) and the `MsDos` function. `LASTDRV.PAS` in Listing 2-5 shows this.

### Listing 2-5: LASTDRV.PAS

```
{ LASTDRV.PAS -- uses only documented DOS }
program LastDrv;
uses Dos;
var
  r = Registers,
  lastdrive = word;
begin
  with r do begin
    ah = $19;           { Get Current Disk }
    MsDos(r);
    dl = al;
    ah = $D;           { Select Disk }
    MsDos(r);
    lastdrive := al;
  end;
  WriteLn('LASTDRIVE=', Chr(Ord('A') - 1 + lastdrive));
  Halt(lastdrive);
end
```

Note that Pascal's *with* construct allows us, for example, to refer to fields of the `Registers` record as *ah* rather than *r.ah*.

The compiled-in version of Turbo Pascal can turn `LASTDRV.PAS` can into `LASTDRV.EXE`.

`tpc lastdrv.pas`

Even so, Turbo Pascal's reputation for producing extremely tight code, the resulting Turbo Pascal executable file is only 2K. The smallest C version is about 4K.

There is one caveat with respect to using `LASTDRV` with making DOS calls from Turbo Pascal for Windows (TPW). For example, instead of using the DOS unit, you must use the `WinDos` unit. See Chapter 3 on a more general discussion of calling DOS from Windows programs. For TPW specifically, see Rick Koenig's excellent book *PC Managers: Turbo Pascal for Windows Techniques and Tricks*, discussed in this subject at length.

### DOS Calls from BASIC

Finally, what about BASIC? `LASTDRV.BAS`, the version of `LASTDRV` in Listing 2-6, displays the `LASTDRIVE` letter, and returns the numeric value of `LASTDRIVE` to the DOS `ERRORLEVEL`.

### Listing 2-6: LASTDRV.BAS

```
REM LASTDRV.BAS -- uses only documented DOS
REM $INCLUDE: 'db.b1'
SUB DOSEEXIT(errorlevel)
  CLOSE
  DIM Regs AS RegType
  Regs.ax = BH4CD0 + errorlevel ' Terminate Process
  CALL INTERRUPT($M2), Regs, Regs
  PRINT "this is never executed"
END SUB

DIM Regs AS RegType
Regs.ax = BH1900 ' Get Current Disk
```

```
CALL INTERRUPT($H21, Regs, Regs)
Regs.dx = Regs.ax
Regs.ax = $H0000 * Select Disk
CALL INTERRUPT($H21, Regs, Regs)
!astdrv = Regs.ax AND $HFF
PRINT LASTDRIVE=; CHR$(ASC("A") + 1 + !astdrv)
CALL BOSEXIT(!astdrv)
END
```

To turn this source code into an executable file, you can use either Microsoft QuickBASIC or the Microsoft BASIC 6.0 compiler, using the BASIC 6.0 compiler, the command is

```
bc /o lastdrv.bas,
Link lastdrv,,,qb.lib,
```

If you are using Quick BASIC to produce a stand-alone executable file, the proper execution is

```
qb lastdrv.bas /L qb.lib
```

Using the BC - O switch to produce a stand-alone executable file from within QuickBASIC is mandatory. Surprisingly, as it seems, Microsoft BASIC has no problem for returning control codes to DOS. In order to determine the value of LASTDRIVE with DOS 3.30 (CORRECTED), I ASSEMBLED BAS with the subroutine DOSEALL which directly calls MS-DOS function 40h. I am not familiar with Return Code and some systems directly bypassing BASIC sometimes can remove this does not work from an executable file that uses the BASIC runtime modules, for example, MS-DOS 3.30 Direct calling INT 21h, function 40h, now an executable that uses the BASIC runtime modules can easily hang the machine.

There's a major problem because the system time will return after calling INT 21h function 40h, it does not end the normal BASIC execution and BASIC never gets any clean up after use of the result is that the user is lost when you refer to the DOS program. Thus, although this code shows how to make low-level connections from Microsoft BASIC to any and all programs, present software BASIC has many features waiting for users, programmers, and gatekeepers, and it gives access to the operating system integrators, and not one of them.

MS-DOS 3.0 and higher come with QBasic. While an excellent addition to DOS in many ways, QBasic doesn't match basic support the CALL INTERRUPT feature that LASTDRIVE uses. The CALL ABSOLUTE POKE and DEL SEG instructions, such as MONEY BASIC's program source files, are not available with MS-DOS, could be used to compile LASTDRIVE, but this method seems worth the effort. And while basic QBasic included with DOS should have made it possible to do many things, but this is a BASIC 6.0, unfortunately, the QBasic RUN command causes too much screen flow, not to be generally useful for creating batch files and files for DOS.

## Using Undocumented DOS

Quarterdeck's expanded memory manager QEMM comes with a program called LASTDRIVE.COM, one of whose uses is to report the value of LASTDRIVE. Interestingly enough, this program does not use the undocumented function 0F1. Instead, it uses undocumented function 52h. We won't see why until later. For now, though, the point is simply that if Quarterdeck can do it, so can you.

What better place to start using an undocumented DOS that with a program like LASTDRIVE whose operation we already know well? In the next section we again show how to create the LASTDRIVE utility in assembly language, C, Fortran, Pascal, and QuickBASIC, but this time using undocumented DOS. In particular, we highlight the larger role that the DOS version number plays when using undocumented DOS.

We start through the process of using a standard DOS programmer's reference as if it were an index, approach of a handbook of mathematical functions, trying to find a tool that would let you write a FATDRIVE file. We eventually find a single function called Get Last Drive, but we did find two other functions called Get Disk and Select Disk that together achieve the same effect.

It is a little surprising that these functions are similar to scandisk, getdisk. But there's something logical about it. Why would DOS return the total number of bytes when you set the current drive? MS-DOS probably keeps this value of LASTDRIVE somewhere internally. Is there some way to find it?

### Disassembling DOS

Of course, you can't disassemble the code for MS-DOS and see how it implements INT 21h function 0Eh, because you don't know it produces the value that is returned by the AX register. Chapter 6 discusses DOS's internal functions, but you can't view them here, so let's look at what actually happens when we call INT 21h function 0Eh.

Figure 2-7 shows the disassembly of the first place that we saw you had to take it to get the value of the current drive. Below is the DOS handout for INT 21h function 0Eh. Microsoft's SYMDEB disassembler, running under DOS 6.0 with DOS 386CPU, produced this disassembly. Microsoft's own disassembler, which is much more expensive, produces a nicer result than SYMDEB.

#### Listing 2-7: SYMDEB Disassembly of INT 21h Function 0Eh in DOS 6.0

```

-> fdc9 4c68      MOV     AL,DL
fdc9 4c68 8AC7      INC     AL
fdc9 4c68 1EFD      CALL   AAB7
fdc9 4c6c 6B85E     CALL   AAB7
fdc9 4c6f 7304      JB     4C75
fdc9 4c71 56A736D3    MOV     SS,EDI361,AL
fdc9 4c75 56A747D0    MOV     AL,SS,EDI471
fdc9 4c79 C3          RET
  
```

This code is a little difficult to follow, so Listing 2-8 contains another version, suitably commented and decorated.

#### Listing 2-8: Commented Disassembly of INT 21h Function 0Eh

```

select_disk_0E proc near
fdc9 4c68      MOV     AL,DL           ; DL = zero-based drive code
fdc9 4c6a      INC     AL             ; make one-based
fdc9 4c6c      CALL   $select_disk   ; call internal function
fdc9 4c6f      JB     fail            ; jump if carry
fdc9 4c71      MOV     DOS_65 [CORRUPTIVE],AL ; 011E 0556
fail:
fdc9 4c75      MOV     AL,DOS_DS [LASTDRIVE] ; 011E 0067
fdc9 4c79      RET
select_disk_0E endp
  
```

The function labeled `select_disk_0E` is, of course, `select_disk_0E`; among other things, it calls the DOS internal set drive function. It is the same piece of code that is invoked if you issue the command `total (INT 21h AX=121E) [DRIVE]`, since the only way to specify DOS can call this function for the sake of setting the INT 21h AX=121E is to call two internal functions to setup the DOS file system (Directory Structure Tables) and call INT 21h AX=121Eh and set drive code INT 21h AX=1217h.

It is a little surprising that `select_disk_0E` sets the carry flag, indicating failure, `select_disk_0E` will always succeed if LASTDRIVE is 1. It is possible that we could pass any illegal drive number into INT 21h to get a carry flag, since we can trick `select_disk_0E` by calling LASTDRIVE. In Listings 2-1 through 2-6, the code handles carry flag function 0Eh to get the current drive, but it is strictly necessary, since all we could get from LASTDRIVE register was 0. Of course, you don't want to pass in an arbitrary *valid*





### Using the Interrupt List

So LASTDRIVE is just a byte stored at offset 21h by the address returned from INT 21h function 52h. We found this information by disassembling DOS. But there's another, easier way to find out this kind of information: from *Ralf Brown's Interrupt List*. When you leaf through Ralf's appendix to his book, you find that DOS internal location for LASTDRIVE is in the middle of SysVars, or the List of Lists. However, it moved around a lot before DOS 4.

**Table 2-1: Movement of LASTDRIVE Within SysVars**

Offset	Size	Description	DOS version
10h	BYTE	number of logical drives in system	DOS 2.x
18h	BYTE	value of LASTDRIVE command in CONFIG.SYS (default 5)	DOS 3.0
21h	BYTE	value of LASTDRIVE command in CONFIG.SYS (default 5)	DOS 3.1-3.4
21h	BYTE	value of LASTDRIVE command in CONFIG.SYS (default 5)	DOS 4.x

SysVars, or the DOS internal variable table, is probably the most important undocumented DOS data structure, and INT 21h function 52h, which returns to the EBX register pair a pointer to SysVars, is the most important and misused DOS function.

Note from Table 2-1 how the offset of the LASTDRIVE field within the DOS internal variable table changed from DOS 3.0 to DOS 3.1. Disassembling just one version of DOS does not, of course, uncover any problem. This sort of undocumented DOS behavior is just what our program has to deal with. What the offset will be in future versions is anyone's guess, and that, of course, is the whole problem with using undocumented DOS.

In future versions of DOS, the LASTDRIVE field might even disappear, breaking whatever program it was dependent on. The only comfort is that, should SysVars change radically, not only will our program quit, but practically all important Microsoft software will break, too.

For example, the DOSMBR sector device driver, crucial to the operation of Windows Enhanced mode, depends on finding LASTDRIVE at SysVars+21h (see Figure 1-10), which is one reason why Windows requires DOS 3.1 or higher and specifically does not run on DOS 3.0. The reliance of key processes of Microsoft software, such as Windows 4.x, on the internal structure of DOS might make this a particularly unlikely to change. However, that perhaps is too much to hope for in a large company, where members of different groups may or may not talk with each other (see the discussion of Chinese Walls in Chapter 1).

Look at the list of the changes to the position of LASTDRIVE within SysVars, and if you look at the appendices for INT 21h function 52h, massive changes throughout SysVars as a whole, one thing becomes clear: constrain INT 21h function 52h, too.

```
INT 21 - DOS 2+ Internal - GET SYSVARS
```

```
AH = 52h
```

```
Return:
```

```
ES:BX -> DOS SysVars
```

From DOS 2.0 onward, this function has been as stable as any documented DOS function. Even standard DOS environments, such as the DOS boxes in OS/2 2.0 and Windows NT, support it (see Chapter 4).

### No Magic Numbers

Because the SysVars structure is so central to DOS programming, many books on the subject end up, as if INT 21h function 52h, somewhere in their sample source code. However, because of their authors' possibly guilty feelings about using undocumented DOS in the first place, sometimes these

books simply leave the code as explained. For example, in the normally well-commented Turbo Pascal source code in an extremely useful book on LAN programming, Craig Chalken's *Blue print of a LAN*, the following code appears without any explanation:

```
regs.ah := $52;
intr($21, regs);
ofs := regs.bx + $22;
seg := regs.es;
while mem[seg:ofs] <> $ffff do
```

Here, the author is using INTR(\$21) function \$21 to get a pointer to SysVars and then using offset 22h in SysVars to get a pointer to the boxed list of device drivers that DOS maintains. Obviously, the author needs to find some device in memory. What's so bad about that? Nothing, except the code appears out of nowhere, with no comment or explanation, like a small piece of magic. The code might as well have been commented "and then a miracle occurs" or even "you are not expected to understand this."

*You are not expected to understand the way UNIX. The canonical comment describing something (page or two) complicated to (rather) explaining properly. From an informal comment in the context searching code of the V6 UNIX kernel. Eric Raymond (ed.), *The New Hacker's Dictionary*, 1991.*

Each uses INTR(\$21) function \$21 but not explain what it does, seems far worse, this using two undocumented call in the first place. As used a word \$21 and 22h, and even 1111h, appear to be far from magic, but they. Let's see if we can't completely demystify INTR(\$21) and \$21.

You can try out this function without even writing a program by using the DOS DEBUG utility. First assemble the DOS call and execute it:

```
C:\WINDOWS>debug
~#
775A:0100 mov ah, 52
775A:0102 int 21
775A:0104 nop
775A:0105
~g 104
AX=5200 BX=0026 CX=0000 DX=0000 SP=FFEE BP=0000 SI=0000 DI=0000
DS=775A ES=028E SS=775A CS=775A IP=0104  NV  JP  EI  PL  NZ  NA  PO  NC
775A:0104 90                                BOP
```

The register dump shows that, in this sample DEBUG session, ES:BX points to 028E:0026:

```
-d es 0026
028E:0020          40 75 8E 02 98 00 8E 02 AC 01          .u. . . . .
028E:0030 70 00 6E 01 70 00 00 02-00 00 49 0C 00 00 EE 0C  p r . p . . . . i . .
028E:0040 00 00 60 0A 00 00 03 05-12 00 F4 09 04 80 99 15          m . . . . .
028E:0050 9F 15 4E 55 4C 20 20 20-20 20 00 90 43 17 8E 02          M U L . . C
028E:0060 47 17 8E 02 47 17 8E 02-43 17 8E 02 43 17 8E 02          G G C C .
```

Aside from the NULL string at offset 0052, it is difficult to find our way around here. But if we hold the SysVars format (see INTR(\$21) in the Appendix) over the DEBUG dump, it makes sense. If we interpret the first four bytes, 40 75 8E 02, as a DWORD pointer, for example, comes out as 028E:7540, a pointer to the first Disk Parameter Block (DPB).

Offset	Size	Description	
00h	DWORD	first Disk Parameter Block	028E:7540
04h	DWORD	list of DOS file tables	028E:0098
08h	DWORD	pointer to C:\IOCHK5 device driver	0070:01A4
0Ch	DWORD	pointer to C:\CDN device driver	0070:016E

Offset	Size	Description	
10h	WORD	max bytes/block	0200
12h	DWORD	first disk buffer	1K 49 0000
16h	DWORD	Current Directory Structures	0K 1E 0000
1Ah	DWORD	pointer to FCB table	0A6D 0000
1Eh	WORD	number of protected FCBs	0000
20h	BYTE	number of block devices	03
21h	BYTE	LASTDRIVE	05
22h	18 BYTES	actual NU1 device driver header	next dev 09F4 0012 attr 8004h strat 1599h intr 159jh name 'NU1'
34h	BYTE	number of JOINed drives	00

We also saw how that evil Turbo Pascal code quoted earlier was adding 22h to the value returned from INT 05h on 52h so that it could get a pointer to the NU1 device, which is the head of DOS's device table. This is one of the most popular uses of INT 21h function 52h, but clearly SOSVARS breaks many other good uses of this byte, it has also been called the DOS List of Lists.

Because of the value of the LASTDRIVE field in SOSVARS, we call LASTDRIVE+1, which is the first entry with a JOINed SYS, does not include a LASTDRIVE statement.

Having seen a little more of what SOSVARS looks like, we can now retrace our steps in building the LASTDRV utility this time, using INT 21h function 52h and the LASTDRIVE field within the DOS internal variable table. You may be thinking that this is a little crazy, because we already know how to get the value of LASTDRIVE using a completely safe and documented function that doesn't change with each new version of DOS. However, we will see later on that using the undocumented entry value of LASTDRIVE can actually be more reliable than using the documented function 0Fh when you're opening up some interesting possibilities. After all, Quarterdeck must have *some* reason for using function 52h rather than function 0Fh.

### Undocumented DOS Calls from Assembly Language

The small assembly language program in Listing 2-9 shows one how to translate the reference material on DOS function 52h in SOSVARS into a working version of LASTDRV.EXE.

#### Listing 2-9: LASTDRV2.ASM

```

; LASTDRV2.ASM -- uses undocumented DOS
        assume cs:_TEXT, ds:_DATA, ss:_STACK
_STACK segment para stack 'STACK'
_STACK ends
_DATA segment word public 'DATA'
msg db 'LASTDRIVE:'
dletter db '?'
db 0ah, 0ah, 'B'
_DATA ends
_TEXT segment word public 'CODE'
public _lstdrv
_lstdrv proc far
        push si
        push bx
        push cx
        mov si, 10h ; assume DOS 3.0

```

```

mov     ax, 3300h      ; Get (genuine) MS-DOS version
xor     bx, bx        ; Don't know if func supported, so zero BX
int     21h
or      bx, bx        ; Is BX still zero, or did it get changed?
jz      do_2130       ; Not supported, so call 21/30
; 21/3306 returns major=BL, minor=BH
; move into AX to make appear as if returned from 21/30
mov     ax, bx
jmp     short got_vers

do_2130:
mov     ax, 3000h     ; Get MS-DOS version number
int     21h           ; major=AL, minor=AH

got_vers:
cmp     al, 2         ; Requires DOS 2+
jl      fail         ; DOS 3+
jne     dos3up       ; DOS 2 x
mov     si, 10h
jmp     short get

dos3up:
cmp     al, 3
jne     ofs21
and     ah, ah        ; DOS 3.0
jz      get

ofs21:
mov     si, 21h      ; DOS 3.1+, DOS 4+

get:
mov     ah, 52h      ; Get SysVars
xor     bx, bx        ; Zero out ES:BX so we can check
mov     es, bx        ; for NULL after INT 21h
int     21h         ; list=ES:BX
mov     cx, es
or      cx, bx        ; Is ES:BX NULL?
jz      fail         ; Function 52h not supported

mov     al, byte ptr es:[bx+si]
xor     ah, ah        ; return LASTDRIVE in AH
jmp     short fini

fail:
xor     ax, ax        ; return 0 in AX
fini:
pop     cx
pop     bx
pop     si

_1stdrv endp

main proc near
mov     ax, DATA
mov     dx, ax
call   _1stdrv
and     ax, ax        ; test for failure
jz      done

mov     bl, al        ; save LASTDRIVE in BL
add     al, 'A'      ; convert LASTDRIVE to drive letter
mov     di, letter, al ; insert into string

mov     ah, 9        ; Display String
mov     dx, offset msg
int     21h

done:
mov     ah, 4Ch      ; Return to DOS
mov     al, bl        ; exit code
int     21h

main endp
_TEXT ends
END main

```

The first assembly routine is being documented (DOS code for displaying output and for exiting to DOS). As the really interesting code and all of the undocumented DOS is in the slightly convoluted assembly routine to be paraphrased in the following pseudocode:

```
offset = 10h,
ver = 0; version 1,
if ver major = 2; return failure,
else if (ver major = 2) offset = 10h,
else if (ver major = 3 or ver minor = 0) offset = 27h,
listOfLists = GetListofLists(),
if (listOfLists == NULL) return failure,
else return listofLists[offset];
```

The goal of the various DOS version number tests is to put the correct location of `EASIDRIVE` into the SI register, so that it can be addressed the base-addressed `SVSWts` that we get back from `DRS` in undocumented function 52. The SI register is preloaded with the offset of `EASIDRIVE` to DOS 3.0, in a (somewhat foolish) attempt to reduce the large number of `JMPs`.

### DOS Versionitis

Not knowing if DOS versions greater than 4.0, we store 27h into the offset. Usually when testing the DOS version, the DOS version is tested for in the *usual old-timey way* (the highest known version is compared to the return of `INT 21h`). I simply don't equate, for example, version 6.0 against null. It's just an application of how to test a constant as DOS 3.0. For example, *PC Magazine* has published several routines that check the DOS version with the `IF` instruction rather than with `JE` or `JZ`. The parameters that are used in these routines are "never" tested "every" one. Microsoft comes out with a new version of DOS. Programmers use their applications, even ones that use undocumented calls, as if were for constant versions. DOS version number. Microsoft introduced the `SETVER` command.

Just as we test DOS versions in the third 3d as one test, obviously we are wanting (but, for example, DOS 6.0 stores `EASIDRIVE` to the same place as DOS 3.3, which in fact it does. When testing version 6.0, the DOS version can be in the test just a portion of you can take the more conservative approach of testing the program under some version of MS-DOS.

Some times it is not clear who the test operation approach is. Absolutely necessary. As an example, the instruction of this book concerned code to access the DOS Swapable Data Area. The code uses `INT 21h` `AX = 0100` and `DOS 3.3` and `INT 21h` `AX = 5100` to check DOS 4.0 and higher. The code checks the computer DOS version. Microsoft released their DOS 3.0 style `AX = 5,000h` at the time. In some cases, it is not clear if a DOS that looks back under DOS 3.0.

The "versionitis" is only a local problem with using undocumented DOS. If your application uses some of the few undocumented functions of the structure, perhaps you should use a routine to test DOS version numbers. On the other hand, there are several double-checks your program might use to ensure that it is not simply overlooking or overlooking some of DOS's arcana, you will see one such double check later in this chapter.

When testing DOS version numbers, remember that DOS function 30h counterintuitively returns the `major` portion of `AX` (the `major` portion of `AX` and the `minor` version number is returned in `AH`). The `major` portion of `AX` is a counterintuitive reminder that a version number such as 3.1 is actually 3.10. In the case of DOS 4.0, the `major` version number in `AH` is neither 01h nor 10h, but 10 decimal 01h.

The most important point about the DOS version number, though, is that it might be wrong. The `SETVER` command allows the DOS version number returned from `INT 21h` function 30h to be set to a value within the operating system. You can see which applications are being taken out by typing `SETVER` at the DOS prompt, for example, in DOS 3.0 and 6.0 `SETVER` tells `WINWORD`, `FEM`, and `EXCEL` that they are real mode DOS stubs of these Windows programs—that they are running under DOS 4.10. It's a sick world out there!

It is important to remember that `SETVER` changes nothing except the version number reported by `INI 21h` function. `SETVER` doesn't change any other aspect of DOS behavior. `WINWORD.EXE` is still running under DOS 6.0, for example, even if `word` is running under DOS 4.10. It is interesting to contrast this with another, backwards-compatible hack of Microsoft's: the undocumented `GetAppCompatFlags` function in Windows. This function changes the actual behavior of Windows on a per-application basis (see *Unconquered Windows*, Chapter 5).

Because `SETVER` just changes the `INI 21h` function, `SETVER` never changes the actual structure of DOS internals. Programs that rely on undocumented DOS need to get the true DOS version number. Fortunately, DOS 5.0 and higher provides a documented way to do this: `INI 21h` function 3306h returns the true DOS version number in `AX`. This function also reports on whether DOS is in ROM and whether DOS-HIGH.

There is one complication here: `INI 21h` function 3306h is supported only in DOS 5.0 and higher. But you don't know whether you are actually running under DOS 5.0 or higher unless you call this function. One of the authors once wrote a piece of code in which he first called `INI 21h` AH=30h and then, if it reported the version number 8.0 or greater, he called `INI 21h` AH=3306h. This was incredible, except because the whole point is that `INI 21h` AH=30h may have been `SETVER`d to occasionally report some version number less than 5.0. But this means you cannot call `INI 21h` AH=3306h without knowing first whether the function is actually supported. And you can't rely on the compatibility being set if you call an unsupported `INI 21h` function.

As seen in Example 2-9, `ANSI DRV 2` ASM and 2.10, `ANSI DRV 4`, the solution is to preadd `0x` with zero before calling `INI 21h` AH=3306h. If the function isn't supported, `0x` will be zero after running the `INI 21h`. If the function is supported, `0x` will be changed to the real DOS version number.

Who would have thought that something as simple as getting the DOS version number would be so complicated? But again, it gets worse. The DOS program `OS 2.2` reports a DOS version number of 204 on OS 2.2 and a 2040 on OS 2.10. Yet, as we saw in Chapter 4, this version number generally behaves like DOS 5.0.

Finally, DR DOS 5.0 and 6.0 report the DOS version number as `5.3`. Chapter 4 shows that there are undocumented DR DOS conventions that return an actual DR DOS version code, rather than the simulated DOS version number.

The reason here is that you should do whatever you can to minimize your dependency on specific versions of DOS because the DOS version number is a complete mess.

## Accessing SysVars

At any rate, as you can see in the following chunk of code, since `SETVER` does not appear appropriate to the version of DOS the program is running under, the rest is `asm`:

```
mov ah, 52h
int 21h
mov al, byte ptr es:[bx+5]
xor ah, ah
```

Actually, and `ANSI DRV 2` ASM, the code is slightly more complicated than this because we have taken the precaution of verifying that undocumented `INI 21h` function 52h is really supported by checking that the pointer to `ENV` is not `NULL`. The `ENV` register pair is loaded with `NULL` prior to invoking `INI 21h` so that if it returns success, simulated DOS environment that doesn't support this function, `ENV` will instead load a reasonable value we can test for:

```
mov ah, 52h
xor bx, bx
mov es, bx
int 21h
mov cx, es
or cx, bx
jz fail
```

Note, however, that the code doesn't check whether INI 71h AH=52h set the carry flag CF. Unless the documentation specifically says that a function sets or clears CF, the state of CF is undefined. The entry for INI 71h function 52h in the appendix to this book says nothing about CF. Thus, far from being an example of careful programmer checking, CF in fact would be a perfect example of relying on *undefined* behavior. Using an undocumented DOS is completely different from relying on undefined behavior.

Although this version of LASTDRIV looks completely different from the version that used only undocumented DOS calls, the result is similar. This version also displays and returns the value of LASTDRIVE. The difference is that now you're getting information straight from the horse's mouth by examining the DOS internal variable table.

### Undocumented DOS Calls from C

This book has spent so much time on the LASTDRIV utility and on various ways of performing DOS calls that you could think there would be nothing more to say. However, the following version, LASTDRV4.C (Listing 2-10), introduces a number of important topics, including the use of far pointers (FAR), the MK\_FP macro for using the DOS version number, and using int86h rather than int87h.

#### Listing 2-10: LASTDRV4.C

```

/* LASTDRV4.C -- uses undocumented DOS */
#include <stdlib.h>
#include <stdio.h>
#include <dos.h>
#define REG_1P
#define REG_FP(seg,ofs) \
    ((void far *)(((unsigned long)(seg) << 16) | (ofs)))
main()
{
    union REGS r;
    struct SREGS s;
    unsigned char far *sysvars;
    unsigned lastdrv_ofs;
    unsigned lastdrv;

    /* Try to get DOS version from 21/5306 call first, since
    this returns the true DOS version number. The 21/30 call
    used to set _osmajor, _osminor can be changed on a
    per-app basis with the DOS SERVER command */
    r.x.ax = 0x5306;
    r.x.bx = 0;
    int86(0x21, &r, &r);
    /* We use 21/5306 actually supported. This is tricky because
    the behavior of unsupported calls is undefined */
    if (r.x.bx == 0) /* BX has changed: func must be supported
    {
        _osmajor = r.h.bl;
        _osminor = r.h.bh;
        printf("21-5306 returns u %02u\n", _osmajor, _osminor);
    }

    /* Get offset for LASTDRIVE within SysVars */
    if (_osmajor < 2) return 0;
    else if (_osmajor == 2) lastdrv_ofs = 0x10;
    else if (_osmajor == 3 && _osminor == 0) lastdrv_ofs = 0x1b;
    else lastdrv_ofs = 0x21;

    /* Get DOS Lists of Lists */
    r.h.sh = 0x52;
    segread(&s);

```



```

s.es = r x bx = 0;
int86x(0x21, &sp, &ip, &si);
/* make sure Function 52h is supported */
if (! s.es && ! r.x.bx)
    return 0;
sysvars = (unsigned char far *) MK_FP(s.es, r.x.bx);
/* Get LASTDRIVE number */
lastdrv = sysvars[lastdrv_ofs];
/* OS/2 DOS compatibility box sets LASTDRIVE to FFh */
if (lastdrv == 0xFF)
    return 0;
/* Print LASTDRIVE letter */
fputs("LASTDRIVE=", stdout);
putchar('A' - 1 + lastdrv);
putchar('\n');
/* return LASTDRIVE number to DOS */
return lastdrv;
}

```

If you compare `LASTDRIVE.C` with the earlier versions that used only downward DOS calls you will notice a number of significant differences. Rather than call `INT_21` function 30h to get the DOS version number, as the assembly-language version did, `LASTDRIVE.C` now uses the global variables `osmajor` and `osminor` provided by most C compilers for the PC. (Microsoft C, Watcom 4.38b, and MetaWare High C 386-NTDIBEH declares these variables; Borland C++ 3.51 declares them.) It is important to remember that in DOS 3.5, for example, `osmajor` is 30, `osminor` is 3, and not 3 and not 0x30. Just note that even `osmajor` and `osminor` come with call `INT_21`h function 3306h. (You're running under a version of DOS that supports this function, use its return value to possibly change `osmajor` and `osminor`.)

Because DOS function 52h returns the address of `sysvars` in `ES:IP`, and because `in86x` doesn't handle segment registers such as `ES`, you need to use `in86x` and `struct SREGS`. You'll need to pass any segment registers *into* Function 52h, so `seg` is as long as does `char`. (Whatever `struct SREGS` holds before calling `in86x`. Note the `seg` is a good idea to call the `seg` and `func` to load the `struct SREGS`, as this example does, because `convex` tries to save your code to a protected mode DOS, extended it will be changed, but can never load the segment registers with garbage values, even if these registers are security not used.)

Because `sysvars` is part of DOS and not located inside your program, it must be addressed with a four-byte far pointer. The `l` variable does just this, what to find this address, and is declared as `char far * lastdrv` as in `lastdrv.c`. This allows us to peek it and poke DOS's global variable, truly even from a C program that otherwise uses only two-byte near pointers.

After DOS function 52h has returned the address of `sysvars` in `ES:IP`, `in86x` returns it in `eax` and `ecx`. How do you move these into `lastdrv`? `LASTDRIVE.C` uses the macro `MK_FP`, which, as its name implies, makes a far pointer from a segment and an offset. This far pointer is provided in the DOS header file with Borland C++ 3.51, but unfortunately, not with Microsoft C. `LASTDRIVE.C` uses the C preprocessor to define a `MK_FP` macro if one is not already present. When the definition of `MK_FP` makes it appear as if it is performing a shift left 16, now good C compiler for the PC turns this code:

```

void far *fp = ((void far *)(((unsigned long)(seg) << 16) | (ofs)))
into this
mov ax, _seg
mov dx, _ofs
mov word ptr _fp, dx
mov word ptr _fp+2, ax

```

You can vary the output of compilers by compiling with the `-Ea` or `-Ea` switch in Microsoft C, for example, or the `-N` switch in Borland C++.

But besides using the `MK_FP` macros in Microsoft C, you could also use the following construct:

```
FP_SEG(doslist) = a.seg;
FP_OFF(doslist) = p.x.bq;
```

`FP_SEG` and `FP_OFF` are two other important macros for PC systems programming in C. Whereas `MK_FP` constructs a far pointer from a segment and an offset, `FP_SEG` and `FP_OFF` extract the opposite operation. `FP_SEG` extracts the segment of a far pointer, and `FP_OFF` extracts the offset. Microsoft's version of `FP_SEG` and `FP_OFF` are C values and can therefore be assigned to:

The `SYSTEM` version of `LASTDRIVE` also does a bit more work than the assembly language version. In the `SYSTEM` version of `LASTDRIVE` (and `LASTDRIVE.C` checks to see if `LASTDRIVE` is 011h, which is the value of the DOS 2.10 BIOS compatibility box, also known as the parity box, used for the `LASTDRIVE` entry in its version of `SysVars`. A program using this compatibility box thinks it's running under DOS 10.00, so you might think the program should simply fail if `osversion < 10`.

However, the support for undocumented DOS was improved in each version of the DOS box, so there were some cases that you could find in a system that was from this simulated DOS environment. For instance, the DOS boxes 2.05, 2.10, which might ask as DOS version 2.00, do provide proper support for `LASTDRIVE` and for most other fields in `SysVars`, as well. It is worth noting that, although the DOS boxes are not identical to the double disks, the DOS 2 compatibility box closely resembles DOS 5.0 with `SHARE` and `cached`. See Chapter 4 for details of DOS simulation in DOS 2.

**What, No Structures?**

For most of this entry, the big question in `LASTDRIVE.C` is "Where are the structures?" You need only look at the entry for DOS function 52h and `SysVars` in the appendix to see that all these offsets were to a structure represented with a C structure. You might even ask why this book doesn't present an `UNDOC.H` include file.

The reason we do not have an `UNDOC.H` include file for you is that programs that use undocumented DOS entry points should use only a few of them. An `UNDOC.H` file could be an invitation to write some undocumented DOS calls. We don't want to promote undocumented DOS as yet another application programming interface (API) consisting of several hundred new functions and data structures.

There is an additional more serious problem with using data structures in undocumented DOS programming: this becomes clear as we discuss the next program, `LASTDRIVE.C` (Listing 2.11), which uses C structures to represent much of `SysVars`.

**Listing 2.11. LASTDRVS.C**

```
/* LASTDRVS.C */
#include <stdlib.h>
#include <stdio.h>
#include <dos.h>
#ifndef MK_FP
#define MK_FP(seg,ofs) \
    ((void far *)(((unsigned long)(seg) << 16) | (ofs)))
#endif
#pragma pack(1)
#define SYSVARS_DECB 12
typedef struct t
    unsigned short retrycount, short retrydelay;
    void far *currdiskbuff;
```

```

void near *unreadcon;
unsigned wcb;
void far *dcb, far *filetable;
void far *clock, far *con;
union {
    struct {
        unsigned char numdrive;
        unsigned wcb;
        void far *first_diskbuff;
        unsigned char nul[18];
    } dos2;
    struct {
        unsigned char numblkdev;
        unsigned wcb;
        void far *first_diskbuff, far *curndir;
        unsigned char lastdrive;
        void far *stringarea;
        unsigned size_stringarea;
        void far *fcbtab;
        unsigned fcb_y;
        unsigned char nul[18];
    } dos30;
    struct {
        unsigned wcb;
        void far *diskbuff, far *curndir, far *fcb;
        unsigned numprotfcb;
        unsigned char numblkdev, lastdrive;
        unsigned char nul[18];
        unsigned numjoin;
    } dos31; /* and higher */
} vers;
} SysVars;

main()
{
    union REGS r;
    struct SREGS s;
    SysVars far *sysvars;
    unsigned lastdrive;
    /* Try 21/3306 first */
    r.x.es = 0x3306;
    r.x.bx = 0;
    int86(0x21, &r, &r);
    if (r.x.bx != 0)
    {
        _osmajor = r.h.bl;
        _osminor = r.h.bh;
    }

    /* No SysVars in DOS 1.x */
    if (_osmajor < 2) return 0;

    /* Get SysVars */
    r.h.ah = 0x52;
    segread(&s);
    s.es = r.x.bx;
    intdosx(&r, &r, &s);
    if (!s.es && r.x.bx)
        return 0;
    sysvars = (SysVars far *) NK_FP(s.es, r.x.bx - SYSVARS_DECR);

    /* Get LASTDRIVE value, depending on DOS version */
    if (_osmajor == 3 && _osminor == 0)
        lastdrive = sysvars->vers.dos30.lastdrive;
    else if (_osmajor == 2)
        lastdrive = sysvars->vers.dos2.numdrive;
    else

```

```

    lastdrive = sysvars->vars dos31 lastdrive,
    /* print LASTDRIVE letter, return LASTDRIVE number */
    printf("LASTDRIVE=%c\n", 'A' - 1 + lastdrive);
    return lastdrive;
}

```

If you look up, over strict SysVars, you should understand why this is sometimes called the List of Lists. Most of the fields are just pointers to other data structures, including the list of DOS Memory Counts Blocks, the list of Drive Parameter Blocks, the DOS device chain, and the File Control Block. Furthermore, *not* a complete SysVars structure; those other fields would each use, for example, "H B bar" or "DPR bar" rather than "vnd bar".

The strict SysVars uses a C union to manage the differences between DOS versions. Unions help represent the changes that each version of DOS brought to SysVars. Each component of a C union is allocated storage starting at the beginning of the union, and the size of a union is the amount of storage necessary to represent its largest component. In other words, as in a variant record in Pascal, the components are overlaid. In this union version of the strict SysVars, the same block of memory can be viewed as a struct dos2, a struct dos40, or a struct dos31.

The `#pragma pack 1` is essential. By default, C compilers for the PC align structures on word (two byte) boundaries. For strict structures to correspond exactly with the layout of the DOS internal variable table, you need to pack the structure on byte boundaries. Otherwise, an unaligned char followed by an unaligned short would occupy four bytes, not three, and the structure would not reflect DOS internal variable table.

Note that the program creates, not a strict SysVars, but a *fat pointer* to a strict SysVars. The memory for the structure already exists inside DOS. Also, to get at some of the useful variables located just below the pointer returned from `IN 1 21b`, `haction`, `s2b`, the program uses `ES:BX:12` rather than `ES:BX`.

This example demonstrates a fundamental problem with using data structures when working with undocumented DOS structures: they are unfixable. The C compiler, seeing a reference such as `doslist->vars->dos31->lastdrive`, simply turns this reference into an offset into `doslist`. The compiler computes these offsets. The program itself does not compute the offsets while it is running, so the offsets don't change at run time, based on conditions such as different versions of an operating system.

You could use some of the sequence-numbering and tagging features of `dos` to create a ListOfLists structure that responded to the DOS version number. When working with undocumented DOS data structures, programmers often wish to do exactly two things. One benefit of C is its ability to implement exactly what the skip by creating classes that manage and hide the complexity of underlying structures.

When using only one or two fields from an undocumented DOS data structure, however, and when placement of the fields within the structure differs from one DOS version to the next, it is best not to use the structure as a fat pointer, compute offsets, instead. Structures may be self-documenting, but the `vars->dos31->Zs` is not the convoluted expression in Listing 2-11 to extract the LASTDRIVE field from the appropriate component of the union version struct ListOfLists. Note how much simpler it is when you use offsets.

```

if (_osmajor == 3 && _osminor == 0) lastdrv_ofs = 0x18,
else if (_osmajor == 2)         lastdrv_ofs = 0x10,
else                             lastdrv_ofs = 0x21,
lastdrv = doslist[lastdrv_ofs];

or

lastdrv_ofs = (_osmajor == 3 && _osminor == 0) * 0x18
             (_osmajor == 2) * 0x10
             /* otherwise */    0x21 ;

lastdrv = doslist[lastdrv_ofs];

```

or the even more compact C expression which also uses the C ternary conditional operator, in the next version of this utility, LASTDRIVE.C in Listing 2-12.

### Listing 2-12: LASTDRIVE.C

```

/* LASTDRIVE.C */
#include <stdlib.h>
#include <stdio.h>
#include <dos.h>
#ifdef __TURBOC__
#define ASM asm
#else defined(_MSC_VER) && (_MSC_VER > 600)
#define ASM _asm
#endif
#error Requires inline assembler
#endif

unsigned _dos_lastdrive(void)
{
    unsigned char far *sysvars;

    /* Try 21/3306 first */
    ASM mov ax, 3306h
    ASM xor bx, bx
    ASM int 21h
    ASM or bx, bx
    ASM jz osmajmin_ok
    ASM mov byte ptr _osmajor, bl
    ASM mov byte ptr _osminor, bh
osmajmin_ok:
    if (_osmajor < 2)
        return 0;
    ASM mov ah, 52h
    ASM int 21h
    ASM mov word ptr sysvars+2, ax
    ASM mov word ptr sysvars, bx
    return sysvars[(_osmajor == 3 && _osminor == 0) ? 0x16 :
        (_osmajor == 2) ? 0x10 :
        /* otherwise */ 0x21];
}

main()
{
    unsigned lastdrive = _dos_lastdrive();
    if (lastdrive == 0xFF)
        return 0;
    fputs("LASTDRIVE=", stdout);
    putchar('A' + lastdrive);
    putchar('\n');
    return lastdrive;
}

```

The other item of interest in LASTDRIVE.C is the use of inline assembly within the function `_dos_lastdrive`. The inline assembly language often seems like an invitation to produce extremely ugly code. C programmers encountering inline assembly language for the first time want to forget all about subroutines. Especially when working with the combination of undocumented DOS and inline assembler you should remember to use subroutines. But also remember our very first rule: to preserve the DI, SI, DS, SS, and SP registers. The inline assembler in `_dos_lastdrive` only changes AX, BX, and ES so you're okay here. The name was chosen to conform to the Microsoft C naming convention: `_dos_getdrive`, `_dos_setdrive`, etc.

**Undocumented DOS Calls from Turbo Pascal**

Turbo Pascal programs that make undocumented DOS calls are similar to those that make documented calls, except that—as with assembly language and C—such programs need to be especially aware of the version of MS-DOS under which they are running. The following program, `LASTDRV2.PAS` (Listing 2-13), uses the function `DosVersion` (added in Turbo Pascal 5.0

**Listing 2-13: LASTDRV2.PAS**

```

{ LASTDRV2.PAS }
program LastDrv,
uses Dos,
var
  r registers,
  lastdrv_ofs Word,
  lastdrive Word,
  vers Word,
begin
  { determine offset of LASTDRIVE within SysVars }
  lastdrv_ofs = $21,
  vers = DosVersion,
  case Ln vers of
    0 : Halt(0); { DOS 1 }
    2 : lastdrv_ofs = $10,
      $ { if Hi(vers) = 0 then lastdrv_ofs := $1B;
  end,
  { Get pointer to SysVars }
  with r do begin
    ah = $52,
    es = 0, bx = 0,
    MemDos(r),
    if es = 0 and ebx = 0 then
      Halt(0),
      lastdrive = Mem(es:bx+lastdrv_ofs),
  end,
  if lastdrive = $FF then
    Halt(0),
  { Print LASTDRIVE letter, return LASTDRIVE value }
  WriteLn('LASTDRIVE ', Chr(Ord('A') + lastdrive)),
  Halt(lastdrive),
end

```

If you are working with a version of Turbo Pascal earlier than 5.0 and don't have the `DosVersion` function, it is easy to write your own:

```

function DosVersion : Word;
var
  r registers,
begin
  with r do begin
    ah := $3000; {should call $5306 too!}
    MemDos(r),
    DosVersion := ah;
  end,
end,

```

Note that `LASTDRV2.PAS` uses the predefined Turbo Pascal array `Mem[]` (to peek at `SesVars`, `Mem[]`, `Mem[]`, and `Mem[]`) map onto the first megabyte of physical memory in the machine, you access them with a segment offset index such as `Mem[seg:ofs]`.

Rather than peek at a raw physical memory address with Mem [ ], you could use a data structure. Just as structures and unions can be used when making undocumented DOS calls from C, so variant records can be used from Turbo Pascal, as shown in LASTDRV3.PAS (Listing 2-14).

#### Listing 2-14: LASTDRV3.PAS

```

( LASTDRV3.PAS )
program LastDrv,
uses Dos,
type
  Dos20 = record
    numdrives : Byte;
    maxbytes : Word;
    first_diskbuff : Longint;
    nul : array [1..18] of Byte;
  end,
  Dos30 = record
    numblkdev : Byte;
    maxbytes : Word;
    first_diskbuff : Longint;
    currdtr : Longint;
    lastdrive : Byte;
    stringerea : Longint;
    size stringerea : Word;
    fcbtab : Longint;
    fcb_y : Word;
    nul : array [1..18] of Byte;
  end;
  Dos31 = record { DOS 3.1 and higher }
    maxbytes : Word;
    diskbuff : Longint;
    currdtr : Longint;
    fcb : Longint;
    numprotcb : Word;
    numblkdev : Byte;
    lastdrive : Byte;
    nul : array [1..18] of Byte;
    numjoin : Word;
  end,
  ListOfLists = record
    shareretrycount : Word;
    shareretrydelay : Word;
    currdiskbuf : Longint;
    unreadcon : Word;
    mcb : Word;
    dpb : Longint;
    filetable : Longint;
    clock : Longint;
    con : Longint;
    case Word of
      20 : (dos20 : Dos20);
      30 : (dos30 : Dos30);
      31 : (dos31 : Dos31);
    end,
var
  lastdrive : Word;
function GetLastDrive : Word;
var
  doslist : ^ListOfLists;
  r : registers;
  vers : Word;

```

```

begin
  ( Get pointer to SysVars )
  with n do begin
    ah $32,
    es 0, bx 0,
    Radix:=r,
    if (es = 0) and (bx = 0) then begin
      GetLastDrive := 0;
      Exit,
    end,
    doslist := PtrOf( es, bx - 32);
  end,
  ( _LASTDRIVE offset depends on DOS version )
  GetLastDrive := doslist^ dos31 lastdrive,
  vers := Drivers on
  case of vers: not
  0 : GetLastDrive := 0; ( DOS 1 )
  2 : GetLastDrive := doslist^ dos20 numdrives;
  3 : if $ifvers = 0 then
      GetLastDrive := doslist^ dos30 lastdrive,
  end,
end,
begin
  lastdrive := GetLastDrive;
  if lastdrive = 0 then
    Halt(0);
  Write n 'LASTDRIVE ', (chrOrd('A' - 1) + lastdrive);
  Halt(lastdrive);
end.

```

THE DOS2PASCAL has used a few interesting structures, and doesn't adjust itself to the DOS version (as the newer DOS2PASCAL2.PAS which simply uses common offsets). This is the same trick of course as we use in our own code, the 4 programming languages. One solution is to use the object-oriented features added to Turbo Pascal, and to use the new SysVars, New Registers, IV, Main, etc. Turbo Pascal 6.0 *to improve readability* contains several deeply documented structures, I'll use a few version-sensitive objects that map DOS internal data structures.

As you can see, the DOS2PASCAL2.PAS is a little more complete, under Turbo Pascal for Windows. Naturally, structures and constants (DOS calls) can be automatically generated. The next chapter contains an in-depth discussion of such a tool, the DOS2PASCAL2.PAS, which can access the DOS internal data structures from professional mode Windows programs, this focuses on C. For an introduction related to TPW, see the chapter "Access: Real Mode" or Real Registers IV, Main, etc. Turbo Pascal to Windows: Techniques and Utilities.

### Undocumented DOS Calls from BASIC

In the BASIC version of the LASTDRIVE program, which uses the undocumented DOS calls, I received a message: "All statements are under control, next code" (MS-DOS 1.1). The second BASIC version of LASTDRIVE (listing 2.15) uses the undocumented SysVars structure, this version requires a DOS 6.0 or higher, this is the only structure that is the object of LASTDRIVE within SysVars.

#### Listing 2.15 LASTDRV2.BAS

```

REM LASTDRV2.BAS
REM $INC=code 'OR BIT'
DEF FNHL x = x \ $H100
DEF FNLO 'x = x AND $FFF
FUNCTION DOSVERSION
  DIM Regs AS RegType
  Regs.es = $H3000 'should call $H3306 too'

```



```

CALL INTERRUPT(&H21, Regs, Regs)
DOSVERSION = Regs.ax
END FUNCTION

SUB DOSEXIT(errorlevel)
CLOSE
DIM Regs AS RegType
Regs.ax = &H400 + errorlevel
CALL INTERRUPT(&H21, Regs, Regs)
END SUB

REM based on DOS version number, find offset of LASTDRIVE
lastdrvofs = &H21
vars = DOSVERSION
IF FNLO(vars) <= 3 THEN DOSEXIT(0)
IF (FNHI(vars) = 3) AND (FNHI(vars) = 0) THEN lastdrvofs = &H1B

REM get address of SysVars
DIM Regs AS RegType
Regs.ax = &H5200
Regs.es = 0
Regs.bx = 0

REM to use current value of DS, set to -1
Regs.ds = -1
CALL INTERRUPTX(&H21, Regs, Regs)
IF (Regs.es = 0) AND (Regs.bx = 0) THEN DOSEXIT(0)

REM peek at LASTDRIVE field within SysVars
DEF SEG = Regs.es
lastdrv = PEEK(Regs.bx + lastdrvofs)
IF lastdrv = &HFF THEN DOSEXIT(0)

REM print LASTDRIVE letter, return LASTDRIVE number
PRINT "LASTDRIVE = ", CHR$(ASC("A") + lastdrv)
CALL DOSEXIT(lastdrv)
END

```

Once INT 21h function 21h returns the address of SysVars, the EBX register of LASTDRIVE2.BAS uses DEF SEG and PEEK to read the LASTDRIVE field. The call to DOSCALLS from BASIC includes instructions for computing this value as a string variable. As noted there, the version of QBASIC included with MS-DOS unfortunately does not support CALL INTERRUPT.

## When Not To Use Undocumented Features

The last few sections described in detail several ways of using undocumented features that may not be available. It is advisable to use MS-DOS's well-documented function interface. This could be compared to an American who learns Japanese and then uses a newly acquired skill only to watch Arrians movies dubbed into Japanese.

This provides us with a few examples of when *not* to use undocumented DOS. If there is a way to perform an operation using the documented DOS program's interface, use it. *One of our goals* is to use the documented interfaces. If there is a seemingly convenient way to accomplish some task using the undocumented calls described in this book, and a conventional way involving documented calls, use the documented calls. You'll see a good example of this in Chapter 7 where we discuss a brief temptation to use INT 29h.

The "Microsoft's best" approach to programming is the desire to use a function simply because it is there. This is when you are experimenting with a new operating system, but it is not a good commercial software. One of our worries as producers of this book was that a single customer, if they ever use an undocumented DOS, please don't use undocumented DOS when documented DOS will do. Using undocumented calls and data structures ties your program to a particular implementation of the operating system, making it harder for the operating system to change, and making compor-

with the DOS more difficult. The more programs depend on undocumented DOS, the more the whole industry will have to suffer under the backward compatibility strangling hand of MS-DOS.

Having said all this, though, let's remember that lots of successful commercial software for the PC has been generated by DOS features. Certain things can't be done using only the documented interfaces, but you should at a minimum try to do the situation with direct hardware access. Clearly, you should use direct hardware access only as a last resort, yet almost all successful PC software does some direct hardware access. A lot of DOS programming takes place in this area of "last resort."

## Verifying Undocumented DOS

Actually, there's one good reason for using undocumented DOS even when there is equivalent documented functionality. It would be nice to have a way to perform a baseline validation of the usability of an undocumented DOS in any given environment. Obviously, the best way to validate a value computed using undocumented DOS is to compare it to a known value with which it should be equivalent.

For example, we can't help but check the results of undocumented DOS against documented DOS because, if we can't do so, we'd be using undocumented DOS in the first place! To check that `dos_getdrive` really does return sectors per disk or sectors per cylinder (as you might expect), we can compare the value returned by `dos_getdrive` from the `SYSTEM` file to the value returned by `dos_getdrive` from the `SYSTEM` file. Then, successfully comparing documented DOS sectors per disk would indicate that something was very wrong.

Of course, you might want your programs to incorporate something similar to the following function in `dos_ohk.c` (Listing 2.16):

### Listing 2.16: `undoc_dos_okay()`

```

BOOL undoc_dos_okay(void)
{
    unsigned char far *sysvars,
    unsigned lastdrv_doc;

    /* could do 21/3306, but if DOS version number from 21/30 doesn't
    accurately reflect genuine DOS version number, then test should
    fail, and undoc_dos_okay() should return FALSE */

    /* get offset for _ASTORIVE within DOS list of lists */
    unsigned lastdrv_ofs = 0x21,
    if (_osmajor >= 5) lastdrv_ofs = 0x10,
    else if ((_osmajor < 5) && (_osmajor != 0)) lastdrv_ofs = 0x1b,
    /* Get DOS lists of lists */
    asm mov ah, 52h
    asm xcr bx, bx
    _asm mov es, bx
    asm int 21h
    _asm mov word ptr sysvars, bx
    _asm mov word ptr sysvars+2, es
    if ! sysvars return FALSE;

    /* use documented DOS to verify results */
    #ifdef __TURBOC__
    lastdrv_doc = getdisk(0xFF), /* don't need getdisk() */
    #else
    dos_setdrive(0xFF, &lastdrv_doc),
    #endif
    return (sysvars[lastdrv_ofs] == lastdrv_doc);
}

```

If `undoc_dos_okay` returns `TRUE` on an `MS-DOS 8.0` it is no guarantee that all code that employs undocumented DOS will work. However, if `undoc_dos_okay` returns `FALSE`, there's a good chance that something is wrong with the underlying operating system, and that your code will have to work

around this. For example, a `dos okay` fails to the OS 2.1.10 DOS box, but it works in the vastly improved Minix DOS boxes of OS 2.2.0.

An interesting question is exactly what to test for. The test in Listing 2.16 is probably too liberal: a true `dos okay` would need to test for more than the following indications of `LASTDRIVE`. But one can make a DOS test that is so stringent and arbitrary that only MS-DOS could hope to pass. A good example is the `WARD` code we examined in Chapter 4. It is code is essentially Microsoft's version of a `dos okay`. Another example is Microsoft's `SETSP0` test discussed in Chapter 4.

## Making Modifications

We noted earlier that QDOS checks `LASTDRIVE.COM` directly, uses the undocumented rather than the documented technique for retrieving the value of `LASTDRIVE`. One reason for this seemingly outrageous disregard of the normal rules of good behavior is that `LASTDRIVE.COM` can also be used to *change* the value of `LASTDRIVE`. This is a good example of something that requires undocumented DOS.

To change the value of `LASTDRIVE` is, of course, a simple matter. You've seen where this value is stored in DOS's data segment; changing the value would involve nothing more than writing rather than reading this memory location.

But to *modify* (change) `LASTDRIVE` requires a little more work. It amounts to adding or subtracting entries from DOS's internal drive table. For each logical drive, DOS maintains a `CDN` (Current Directory Structure) As explained in some detail in Chapter 5, see the `LASTDRIVE.COM` sample program, the `CDN` marks the current working directory for each drive in the system. DOS uses these entries of these `CDN` structures. Its pointer to this array is stored in `SYSDRV`, just like `LASTDRIVE`.

`LASTDRIVE` in fact is nothing more than the size of the `CDN` array. Quite often `LASTDRIVE.COM` not only changes the value of `LASTDRIVE` but also releases and creates the `CDN` array. Indeed, the sole purpose of `LASTDRIVE.COM` is to associate the `CDN` structure with memory block 1, which is the sole purpose of `LASTDRIVE` under CONROSS and the original DADSH. `LASTDRIVE.COM` also uses some QDOS conventions such as `FILES.COM` or `UTILS.COM` for creating other DOS internal data structures concerned with INT 21h function 52h.

How does it all work? Of course, to test out would be to disassemble QDOS's `LASTDRIVE` with the 286.COM file. Running Subgroup's `Source` or another `LDOS` source program is a bit awkward, but the results are not so bad. The assembly language code is well written, comments and error messages. However, it is a necessary observation: `LASTDRIVE` breaks the program `INTRSPY` because that early supervisor `LASTDRIVE` calls `INT 21h` function 52h. A few moments' thought about the problem of reworking the `CDN` entries (a bit pseudo-code shown in Listing 2.17).

### Listing 2.17: Pseudocode for Changing `LASTDRIVE`

```
get new_lastdrv from command line,
get sysvars from INT 21h 4h/5h
set LASTDRV_OFS and CDS_PTR_OFS based on DOS version,
old_lastdrv = sysvars[LASTDRV_OFS],
old_cds = sysvars[CDS_PTR_OFS],
sizeof_cds_entry depends on DOS version,
new_cds = allocate(new_lastdrv * sizeof_cds_entry), // 1040h
disable interrupts,
memcpy(new_cds, old_cds, old_lastdrv * sizeof_cds_entry),
mark newly formed CDS entries as invalid drives,
sysvars[LASTDRV_OFS] = new_lastdrv,
sysvars[CDS_PTR_OFS] = new_cds; // do it!
enable interrupts,
```

```
// hope that no one has a dangling invalid pointer to old_cds*
// (in any case, can't free it because don't know how it was allocated)
```

Basically, evaluating the CDS is a matter of allocating a block of memory for the new CDS, copying the old CDS into it, and updating the CDS pointer and LASTDRIVE value in SysVars. It is easy to fit these steps into a working program. Listing 2.18 shows XLASTDRV.C, which for the most part just carries out the actions shown in the pseudocode.

### Listing 2.18. XLASTDRV.C

```
/*
XLASTDRV.C
Andrew Schulman, May 1993
bcc xlastdrv.c ..\chap3\iswin.c

Straightforward clone of GENM_LASTDRIV.COM utility, except
Checks for presence of Windows Enh mode (DOSNGE hangs onto original
CDS pointer from SysVars), refuses to run in DOS box
- Checks for presence of Windows Std mode and task switcher (suggested
  by Ralf Brown). Otherwise, DOS box can be closed, thereby deallocating
  XLASTDRV's CDS*
Checks for presence of DesqView (suggested by Ralf Brown): the new
CDS would be in remapped memory that isn't guaranteed to be around when
a DOS call uses the CDS.
Uses sleazy hack to try to free up previously-created XLASTDRV CDSes
Uses instant TSR technique from Prosize UMDFILES
-- Provides RESTORE command to put back original CDS pointer and LASTDRIVE
  from SysVars, uses sleazy hack entry to save these away. This necessary
  for MSCDEX and other utilities and hold onto original SysVars pointers.
Oh, not such a straightforward clone after all!
*/
```

```
#include <stdlib.h>
#include <stdio.h>
#include <ctype.h>
#include <string.h>
#include <dos.h>
#ifndef __TURBOC__
#include <alloc.h>
#include <dir.h>
#else
#include <direct.h>
#endif

typedef unsigned char BYTE;
typedef unsigned short WORD;
typedef unsigned long DWORD;

#define min(a,b) ((a) < (b) ? (a) : (b))
#ifndef MK_FP
#define MK_FP(seg,ofs) \
  ((void far *)(((DWORD)seg) << 16) | (ofs))
#endif

typedef struct {
  BYTE zero; far *orig_cds, orig_lastdrv, signature[],
  } SLIATY_WACK;

#define GET_SLEAZY_WACK(cds, td) \
  ((SLEAZY_WACK far *) ((cds)+((td)-1)*sizeof_cds_entry))

#define XLASTDRV_SIGN XLASTDRV sleazy hack

// \undoc2\chap3\iswin.c
extern int _is_win(int *pmsj, int *pmsin, int *pmode);
```

```

extern int detect_switcher(void);
void set_genuine_dos_ver(void);
BYTE far *get_sysvars(void);
void set_mcb_owner(WORD para, WORD owner);
void set_mcb_name(WORD para, char far *name);
int get_lastdrive_doc(void);
unsigned desqview(void);
void fail(char *s) { puts(s); exit(1); }

main(int argc, char *argv[])
{
    SLEAZY HACK far *old_hack, far *new_hack,
    BYTE far *sysvars, far *old_cds, far *new_cds,
    unsigned int para,
    int lastdrv_ofs, cdsptr_ofs, new_lastdrv, old_lastdrv,
    int sizeof_cds_entry,
    int do_restore = 0;
    int dummy,

    if (argc < 2)
        fail( XLASTDRV [new value for lastdrv: A-Z] or -restore");

    if (is_win(&dummy, &dummy, &dummy)) // including Std mode (see Listing 3-29)
        fail( Sorry, XLASTDRV can't run under Windows 1,
    if (detect_switcher() != 0) // other than Std mode (see Listing 3-29)
        fail("Sorry, XLASTDRV can't run under a task switcher");
    if (desqview() != 0)
        fail( Sorry, XLASTDRV can't run under Desqview ),

    // get new LASTDRV (or -RESTORE) from command line
    if (strcmp(strupr(argv[1]), "-RESTORE") == 0)
        do_restore = 1;

    /*
    new_lastdrv = 1 + (toupper(argv[1][0]) - 'A'),
    // possibly change osmajor, _osminor, see if ok
    set_genuine_dos_ver();
    if ((_osmajor < 3) || (_osmajor >= 10))
        fail("Unsupported DOS version");
    // get offsets within SysVars
    if ((_osmajor == 3) && (_osminor == 0))
    {
        cdsptr_ofs = 0x17,
        lastdrv_ofs = 0x1b,
    }
    else
    {
        cdsptr_ofs = 0x16,
        lastdrv_ofs = 0x21,
    }

    // old_lastdrv = sysvars[lastdrv_ofs],
    if (! (sysvars = get_sysvars()))
        fail("Can't get SysVars");
    old_lastdrv = sysvars[lastdrv_ofs],

    // see if SysVars looks ok
    // get_lastdrive_doc() is wrong for Novell see Listing 2-19
    if (get_lastdrive_doc() != old_lastdrv)
        fail("Internal DOS data structures wrong"),
    if (old_lastdrv == 0xff)
        fail( LASTDRIVE not supported"),

    // old_cds = sysvars[cdsptr_ofs];
    old_cds = *((BYTE far *) far *) &sysvars[cdsptr_ofs],

```

```

// sizeof_cds_entry depends on DOS version;
if (_osmajor == 3)
    sizeof_cds_entry = 0x51;
else
    sizeof_cds_entry = 0x58,

if (do_restore)
{
    old_hack = GET_SLEAZY_HACK(old_cds, old_lastdrv),
    if (!strcmp(old_hack->signature, XLASTDRV_SIGN) != 0)
        fail("Can't restore original CDS!");
    _disable(),
    _fmemcpy(old_hack->orig_cds, old_cds,
        min(o.d._lastdrv, old_hack->orig_lastdrv) * sizeof_cds_entry);
    *((BYTE far * far *) &sysvars[cdsptr_ofs]) = old_hack->orig_cds;
    sysvars[lastdrv_ofs] = old_hack->orig_lastdrv;
    _enable(),
    printf("Restored CDS to %p (LASTDRIVE=%c)\n",
        old_hack->orig_cds,
        (old_hack->orig_lastdrv - 1) + 'A');
    goto free_old;
}

if (new_lastdrv != old_lastdrv)
    printf("Changing LASTDRIVE from %c to %c\n",
        (old_lastdrv - 1) + 'A',
        (new_lastdrv - 1) + 'A'),

// temporarily bump up new_lastdrv for signature entry
new_lastdrv++,

// new cds = alocate(new_lastdrv * sizeof_cds_entry),
if (!dos_allocate(1+(new_lastdrv * sizeof_cds_entry) >> 4),
    &para, 0)
    fail("insufficient memory"),
new_cds = (BYTE far *) HK_FP(para, 0);

// this taken from Jeff Prosser, UMDFILES (PC MAG, 11/26/91)
// instant TSR! (NEW /D thinks XLASTDRV is TSR)
set_mcb_owner(para, para),
set_mcb_name(para, "XLASTDRV"),

// mark newly-formed CDS entries as invalid drives,
_fmemset(new_cds, 0, new_lastdrv * sizeof_cds_entry);
printf("Moving CDS from %p to %p\n", old_cds, new_cds),
_disable(), // disable interrupts while moving CDS
// copy old CDS entries over to new CDS table
_fmemcpy(new_cds, old_cds,
    sizeof_cds_entry * min(new_lastdrv, old_lastdrv));
// put back actual new LASTDRIVE value requested by user
new_lastdrv--,

// use extra entry for signature, and to keep orig to restore
new_hack = GET_SLEAZY_HACK(new_cds, new_lastdrv),
old_hack = GET_SLEAZY_HACK(old_cds, old_lastdrv);
memcpy(new_hack->signature, XLASTDRV_SIGN);
if (!strcmp(old_hack->signature, XLASTDRV_SIGN) == 0)
{
    // have done XLASTDRV before, copy orig values from old to new
    new_hack->orig_cds = old_hack->orig_cds;
    new_hack->orig_lastdrv = old_hack->orig_lastdrv,
}
else
{
    // this is first time we've run XLASTDRV

```

```

    new_hack->orig_cds = old_cds;
    new_hack->orig_lastdrv = old_lastdrv;
}
// plug new values into SysVars!
*(BYTE far * far *) &sysvars[cdsptr_ofs] = new_cds;
sysvars[lastdrv_ofs] = new_lastdrv;

_enable(), // new CDS in place, re-enable interrupts
// hope that no one has a dangling invalid pointer to old_cds!
// Windows does, MSCDEX does!

free_old:
// if we allocated old one, we can free it
if (!_stricmp(old_hack->signature, XLASTDRV_SIGN) == 0)
{
    if (_dos_freemem(FP_SEG(old_cds)) != 0)
        fprintf(stderr, "Couldn't free up old CDS");
    else
        printf(" Freed old XLASTDRV COS at %p\n", old_cds);
}
return 0;
}

void set_genuine_dos_vers(void)
{
    union REGS r;
    r.x.ax = 0x3306;
    r.x.bx = 0;
    intdos(&r, &r);
    if (r.x.bx != 0)
    {
        _osmajor = r.h.bl;
        _osminor = r.h.bh;
    }
}

BYTE far *get_sysvers(void)
{
    union REGS r;
    struct SREGS s;
    r.h.ah = 0x52;
    s.eax = 0;
    s.ebx = 0;
    intdosx(&r, &r, &s);
    return (BYTE far *) MK_FP(s.es, r.x.bx);
}

void set_acb_owner(WORD para, WORD owner)
{
    *((WORD far *) MK_FP(para-1, 1)) = owner;
}

void set_acb_name(WORD para, char far *name)
{
    _fmemcpy(MK_FP(para-1, 8), name, 8);
}

unsigned desqview(void)
{
    _asm mov ax, 2801h
    _asm mov cx, 4445h /* 'DE' */
    _asm mov dx, 5351h /* 'SQ' */
    _asm int 21h
    _asm cmp al, 0FFh
    _asm je no_desq
    _asm mov ax, bh /* BH=major, BL=minor */
    _asm jmp short done
}

```

```

no_desq
    _osw xor ax, ax    /* return 0 */
done:
    // return value in AX
}

int get_asterive_doc(ward) // should adjust for Novell!
{
    #ifdef __TURBOC__
        return setdisk(0xFF);
    #else
        unsigned lastdrv_doc;
        _dos_setdrive(0xFF, &lastdrv_doc);
        return lastdrv_doc;
    #endif
}

```

After `ALANDRV` runs, `DCS` uses the new CDS it allocates just as if it were `DCS`'s own creation. Sticking the two new values into `SysVars` took care of this. However, if you run `MEM_ID` or the `CDMEM` and review Chapter 7 you will see that these utilities don't know that the CDS has moved. And it isn't possible to tell them either. The `QMEM` `ALANDRV` utility has the same problem. But as it's not these public utilities that accept the new CDS as its own.

To redeflate the CDS to its just-run state `ALANDRV` or `LOADHIGH` `ALANDRV` `ALANDRV` can continue to allocate memory for the new CDS using the Microsoft and Borland `dos` functions, functions which call `INT 21h` to set to 48h. If you run `ALANDRV` under `LOADHIGH` `DCS` will just allocate the memory from a `UMB`. Which would be easy enough to risk. `ALANDRV` `UMB` over-see `IN` `Altair`'s `UMBLETS` entry, for example, this step really isn't necessary. I will explain later how `LOADHIGH` works, see Chapter 7.

With `ALANDRV` `alt` have to worry about is making sure that the memory for the new CDS was a normal data program space. When `ALANDRV` runs, `DCS` walks the Memory Control Block chain and finds out if it was owned by `ALANDRV`. It seems like overkill to make the program a `FSB` so instead of allocating a block of memory for the new CDS, `ALANDRV` changes the block size of a `DCS` `FSB` to free the block when `ALANDRV` exits. In essence, this procedure creates a constant `FSB`. See Chapter 7 for details of the `MCB` chain.

It is essential to free up the old CDS. We can't free the original CDS owned by `DCS`, even though we might use `fs` to specify how it was allocated. But we can free up any earlier CDS that `ALANDRV` had allocated, a mark it as a CDS allocation. `ALANDRV` allocates an extra CDS pointer to the old CDS, and `ALANDRV` starts back. If `ALANDRV` ever finds CDS signature, it can free CDS. Or you can consider the old CDS.

There is one other problem with the software, a key technique of relocating internal `DCS` data structures to `UMB`. Relocation depends on a complex voluntary convention that the operating system will always use `FSB` always through `SysVars`. If a resident program gets a pointer to the CDS, it uses `SysVars` once during initialization and uses it thereafter without referring back to `SysVars`. The operating system, `DCS` or `ALANDRV` or any similar program comes along later and changes the CDS pointer in `SysVars`.

To simplify the Microsoft `CDROM` extensions, `MSCDEX` save away a pointer to the CDS. If `ALANDRV` changes the CDS location, `MSCDEX` can no longer find its `CDROM` drives. `ALANDRV` provides a `RESTORE` switch that puts the CDS back to its original location. Of course, an addition to `fs` will not hurt, as `MSCDEX` works again.

Note, in Listing 2.18 that `ALANDRV` uses `INSTR` `CDROM` Chapter 3 (Listing 3.29) to see if Windows is active. Windows Enhanced mode does not follow this always refer to `SysVars` convention at all. The `DRSMGR` `ASDF` checks `SysVars` once during its initialization and isn't prepared for the `SE` or CDS pointers to change after this. As seen back in Listing 1.10, `DRSMGR` hangs on to `base` as linear addresses, and isn't prepared for `ALANDRV` to change. Thus, while running





Instead of asserting that undocumented DOS is unusable simply because `doslist >lastdrive > setdisk` fails, the improved version of `undoc_dos_okay.c` performs a slightly more complicated test, as shown in Listing 2-19.

### Listing 2-19: OKAY.C

```

/*
OKAY.C -- basic test for validity of undocumented DOS
cl -DTESTING okay.c
*/
#include <stdlib.h>
#include <stdio.h>
#ifdef __TURBOC__
#include <dir.h>
#endif
typedef int BOOL;
#define FALSE 0
#define TRUE (!FALSE)
BOOL network(void),
unsigned lastdrv_network(void),
BOOL undoc_dos_okay(void)
{
    unsigned char far *sysvars;
    unsigned lastdrv_doc;

    /* could do 21/3306, but if DOS version number from 21/30 doesn't
       accurately reflect genuine DOS version number, then test should
       fail, and undoc dos okay() should return FALSE */

    /* get offset for LASTDRIVE within DOS List of Lists */
    unsigned lastdrv_ofs = 0x21;
    if (_osmajor==2) lastdrv_ofs = 0x10;
    else if (_osmajor==3) && (_osmajor==0) lastdrv_ofs = 0x1b;

    /* Get DOS Lists of Lists */
    _asm mov ah, 52h
    _asm xor bx, bx
    _asm mov es, bx
    _asm int 21h
    _asm mov word ptr sysvars, bx
    _asm mov word ptr sysvars+2, es
    if (! sysvars) return FALSE;

    /* use documented DOS to verify results */
#ifdef __TURBOC__
    lastdrv_doc = setdisk(0xFF);
#else
    _dos_setdrive(0xFF, &lastdrv_doc);
#endif
    if (sysvars[lastdrv_ofs] == lastdrv_doc)
        return TRUE;
    else if (network())
    {
        puts("Novell NetWare");
        if (lastdrv_doc != 32)
            puts("NetWare INT 21h DEH looks strange");
        return (sysvars[lastdrv_ofs] == lastdrv_network());
    }
    return FALSE;
}

/* Novell Return Shell Version function (INT 21h AH=EAH AL=01h)
   see INTRLIST on accompanying disk; also see Barry Nance,
   "Networking Programming in C", pp. 117, 341 Z. */

```

```

BOOL network(void)
{
    char buf[40];
    char far *fp = buf,
    _asm push di
    _asm mov ax, 0EAG1h
    _asm mov bx, 0
    _asm .es di, fp
    _asm nt 21h
    _asm xor al, al
    _asm mov ah, bx
    /* If BX still 0, then NetWare not present, return in AX */
    _asm pop di
}

/* Novel. Get number of local drives function (INT 21h AH=0Bh).
See INTRLIST on accompanying disk, or Charles Rose, "Programmer's
Guide to NetWare," p. 731. */
unsigned _astdrv_network(void)
{
    _asm mov ah, 0Bh
    _asm int 21h
    /* AL now holds number of "local drives" (genuine LASTDRIVE) */
    _asm xor ah, ah
    /* unsigned returned in AX */
}

#ifdef TESTING
main()
{
    fputs("Undocumented DOS ", stdout);
    puts( _astdrv_network() ? "ok" : "not ok");
}
#endif

```

Function 0Bh isn't the only DOS function whose behavior Novell NetWare modifies. The NetWare shell inspects every INT 21h function request before DOS itself sees it. The shell decides whether to pass that function request along to DOS or to pass the request over the network to another machine that serves whatever service that DOS machine finally expects the shell to deliver. To pass the INT 21h function request along to DOS it gets to modify any registers before returning control to the application, such as LASTDRIVE, that called INT 21h in the first place. For further details on DOS differs, see under NetWare, see Chapter 4.

## Hooking DOS: Application Wrappers

Seeing that NetWare changes DOS function 0Bh gives us an excuse to look into what it means for a program to "hook DOS." In order to hook DOS, all a program has to do is get the address of the current interrupt handler for INT 21h and then install its own handler for INT 21h. There's nothing difficult or undocumented about this capability. It's built right into DOS itself and is one of the key facilities that makes DOS extensible.

We can simulate NetWare's handling of function 0Bh and provide a realistic example of hooking DOS by using just a few lines of code. Generally programs that hook DOS like Novell's workstation shells are memory resident. However, there is no reason such programs must be ISRs. As FUNC 0F32 in Listing 2-20 shows, it is easier to build a program that instead acts as a shell—a wrapper around another program—see Jim Kyle, "Application Wrappers," *PC Technician*, June 1992. FUNC 0F32 modifies function 0Bh to return 32 just like NetWare, and the cases one of the C spawn functions to run whatever program was specified on its command line. For example,

```
C MUNDOS2\CHAP2>func0e32 lastdrv.exe
LASTDRIVE=E
10 DOS calls
Y changed
```

Any version of `LASTDRIVE` that uses only documented DOS functions is hooked by `FUNC0E32`. For the versions that use undocumented DOS aren't. For example

```
C MUNDOS2\CHAP2>func0e32 lastdrv2.exe
LASTDRIVE=E
10 DOS calls
Y changed
```

`LIST0E32.C` in Listing 2 20 consists of two functions. The function `DOS` is our `INT 21h` handler. Whenever `DOS` is get control each time `anyone` invokes `INT 21h`, `DOS` changes the value that `Func0E32` returns in `AL` to 32. The function also keeps count of how many `INT 21h` calls it has seen and how many it has changed. It is `FUNC0E32` which produces the "10 DOS calls" and "Y changed" output above. In Listing 2 20 main installs `DOS` as the `INT 21h` handler, spawns the program `anyone` from the command line, and then restores the original `INT 21h` handler, which may be `DOS` itself but which, on most PCs, will be some other program that had earlier hooked `DOS`, such as `NetWare`.

### Listing 2 20: FUNC0E32.C

```
/* FUNC0E32.C -- take over INT 21h Function 0EH; return 32 in AL */
#include <stdlib.h>
#include <stdio.h>
#include <process.h>
#include <dos.h>
#pragma pack(1)
typedef struct {
  #ifdef _TURBOC_
    unsigned short bp,di,si,ds,es,ds,cs,bs,es,
  #else
    unsigned short es,ds,di,si,bp,sp,bs,ds,cs,es, /* same as PUSHA */
  #endif
    unsigned short ip,cs,flags,
} REG_PARAMS;

void _interrupt_ for DOS(REG_PARAMS r);
void _interrupt_ for %d)(void);
unsigned long calls = 0;
unsigned long changed = 0;

void fail(char *s) { puts(s); exit(1); }

main(int argc, char *argv[])
{
  if (argc < 2)
    fail("usage: func0e32 [program name] <args ...>");

  old = (void (_interrupt_ for *) (void)) _dos_getvect(0x21),
  _dos_setvect(0x21, DOS); /* hook INT 21h */
  if (!spawnvp(P_WAIT, argv[1], &argv[1]) == 1) /* run command */
    puts("Can't run program!");
  _dos_setvect(0x21, old); /* unhook INT 21h */
  printf("\n%10u DOS calls\n", calls);
  printf("%10u changed\n", changed);
  return 0;
}

void _interrupt_ for DOS(REG_PARAMS r)
{

```

```

_asm pusha
calla++;
if ((r_ax >> 8) == 0x0E) /* If Function DEh. . */
{
    _asm popa
    (*old)(), /* first call old INT 21h handler */
    r_ax = 0x0E20, /* then force #drives to 20h (32) */
    changed++;
}
else /* not for us */
{
    _asm popa
    _chain intrfold(), /* pass to old INT 21h handler */
}
}

```

This code is important not only to illustrate that it is perfectly legitimate (and completely correct) for a company like Novell to change the return value from a DOS function, but also as an example of how to hook a DOS interrupt like INT 21. Some undocumented DOS functions are not for you to call, but on you to *emulate* so that DOS can call them from its own code, if you wish. Such functions are indicated in the appendix with the phrase “Call Emulate rather than Call with.” For example, the DOS network redirector, which incidentally Novell did not so-called recently, is one such set of call-back functions. A network redirector program hooks over INT 21h and looks for calls to API UFs with subfunctions in M.

Listing 2-20 shows you how to hook an interrupt with just a few lines of code. The DOS 44 header files provided by Microsoft Borland, and other PC compiler vendors, declares the functions `dos_getvect` and `dos_setvect`. The `interrupt` keyword defines a function as an interrupt handler. The compiler pushes an image of the CPU registers on a C runtime mode's stack (see PC PARAMS in Listing 2-20). Unfortunatly, you must make the register image as a 286 or 386 with either an 286 and higher processors, whereas Microsoft puts the register image on the 386 or higher. Either way, the interrupt handler can compare the registers image against some value (e.g., 0x0F20).

Of course, the new interrupt handler needs to call back to the previous handler. If you need to call back to get control again after calling the previous handler, the new handler can simply call `dos_setvect` with an expression such as `*old` as the new handler. If you need to get control back to `old`, instead wants to “chain” to the previous handler, it can use the function `chain` that is provided in both Microsoft C 4.53 and as a recent addition in Borland C++ (see `dos_setvect` and `dos_getvect` mentioned check your compiler's runtime library source code, for example `dosrand\src\lib\chainnt.asm`).

FUNCTION 32 is a rather specialized and somewhat esoteric collection of hooks, DOS through an application wrapper. As another example of hooking DOS functions, consider the DOSVER program shown in Listing 2-21. This program takes over INT 21's function `softwarekeys` on back with the DOS version number on an application-by-application basis. It uses the `SETVER` command, but less permanent in its effect.

### Listing 2 21: DOSVER.C

```

/* DOSVER C -- set different DOS version numbers */
#include <stdlib.h>
#include <stdio.h>
#include <process.h>
#include <dos.h>
#pragma pack(1)
void interrupt far *old();
unsigned dosver, old_bx, old_cx;
typedef struct {

```

```
#ifdef __TURBOC__
    unsigned short bp,di,si,ds,es,dx,cs,bx,ax;
#else
    unsigned short es,ds,di,si,bp,sp,bx,dx,cs,ax; /* PUSHA order */
#endif
    unsigned short ip,cs,flags;
} REG_PARAMS;

void interrupt far dos(REG_PARAMS r)
{
    if ((r.ax >> 8) == 0x30)
    {
        r.ax = dosver;
        r.bx = old_bx;
        r.cx = old_cx;
    }
    else
        _chain_intr(old);
}

void fail(char *s) { puts(s); exit(1); }
main(int argc, char *argv[])
{
    int new_major, new_minor;
    unsigned char old_major, old_minor;

    if (argc < 4)
        fail("usage: dosver <major> <minor> <command...>\n"
            "example: dosver 3 31 exe2bin devlod.exe devlod.com");
    if (! (new_major = atoi(argv[1])))
        fail("bad version number");
    if ((new_minor = atoi(argv[2])) < 10) /* e.g. 3.1 to 3.10 */
        new_minor += 10;
    dosver = (new_minor << 8) + new_major;

    _asm mov ax, 3000h
    _asm int 21h
    _asm mov old_major, al
    _asm mov old_minor, ah
    _asm mov old_cx, cx /* OEM, serial# */
    _asm mov old_bx, bx
    if ((new_major != old_major) || (new_minor != old_minor))
        fail("no change");
    printf("Changing DOS version from %u %02u to %u %02u\n",
        old_major, old_minor, new_major, new_minor);

    pid = _dos_getvect(0x21); /* save INT 21h */
    _dos_setvect(0x21, dos); /* hook INT 21h */
    if (spawnvp(P_WAIT, argv[3], argv[3]) == -1) /* run command */
        puts("Can't run program!");
    _dos_setvect(0x21, old); /* unhook INT 21h */
    return 0;
}

```

In addition to its occasional usefulness with programs that are over-conservative in their response to the DOS version number, `DOSVER` and `FUNCBE32` can also be lashed together to thoroughly mess up both `INT 21h` functions `30h` and `0fh` and to see if the `LASTDRV` programs respond properly. For example, the following test runs under `DOS 6.0`, reversely sets the `DOS` version number to `3.31` and then makes function `0fh` return `32`. But `LASTDRV 4` from Listing 2.10 recovers and gets the correct answer.

```
C:\MUNDOC2\CHAP2>dosver 3 31 funcbe32 lastdrv4
Changing DOS version from 6.00 to 3.31
```

```
21/30 returns 3.31
But 21/3306 returns 6.00
LASTDRIVE=#
45 DOS calls
D changed
```

If you haven't written a C interrupt handler before, you may want to experiment with `INTERRUPT32.C` or `DOSVFR.C`, perhaps modifying them to count and display the number of different `INT_21h` calls. Try adding an additional `INT_21h` handler to count the number of `INT_21h` calls an application makes. This should take less than 100 lines of code, yet it would be a serious, scaled-down version of the `INTRSPY` program from Chapter 5.

## On to Protected Mode

Frankly, it is a little difficult to take seriously all the little character-mode DOS programs we've developed in this chapter. In the 1990s, most new development on the PC is being done for Windows, or even for completely different non-DOS operating systems such as OS/2 and Windows NT. Little utilities that print `LASTDRIVE = F` almost seem like a joke. They do something of a "trickiest" "look what I can do with my computer" feel to them.

But issues of MS-DOS systems programming and undocumented DOS rarely haven't gone away. Far from it, Windows applications running in protected mode often need to talk to DOS I/Os and device drivers. In Linux and most other operating systems, there's some way for their Windows applications to get at something from a DOS box (OS/2 and NT emulate DOS, and it's important to know whether any of this undocumented DOS stuff works in new DOS emulators). Undocumented DOS hasn't gone away. As you will see in the next few chapters, it's just become a lot more complicated.





## Undocumented DOS Meets Windows

by Andrew Schulman

The previous chapter has shown how to call undocumented DOS functions and a few—documented—DOS data structures from an assortment of programming languages. But not everyone is as computer- or assembly-oriented as I am, and I can't be sure that all readers will be. Microsoft assigned the MS-DOS operating system to the Intel 8088 microprocessor and, with one of the fastest Pentium or 80486 processors with capabilities of millions. MS-DOS still runs in protected mode, limiting programs to a specific portion of the 1-Gbyte addressable main RAM.

Fortunately, almost no hardware and only a few 8088-based PCs survive, and even a few new developers write real-mode DOS programs. Some Windows applications, managers, bootstraps, and DOS-only development utilities still use MS-DOS real-mode I/O software services. However, for Windows and other protected-mode DOS variants, a little bit of magic in the DOS-to-protected-mode programs can directly access a simple segment of memory on the 80286 and higher processors. The great majority of Windows systems still use real-mode for its access, so that Windows programs run in protected mode.

But even with the new real-mode Windows 3.11 write-protected mode, DOS can be used to access a few Windows-protected-mode services, including unformatted compressed disk files, independent services that use DOS real-mode services, such as DOS AH=0x Windows programs, and real-mode protected-mode of the 80286 and higher microprocessors. This is especially true for segments of extended memory. The programs can indirectly request access to DOS extended BIOS real-mode operating systems. Windows programs can call DOS services from INT 21h or real-mode DOS 3.11 functions from the Windows API, or more likely, a Windows-compatible real-mode API function, which in turn calls INT 21h.

When Windows 3.11 was originally DOS- or program-ready for the implementation to run on a simple INT 21h, the real-mode DOS code, such as Windows 3.11, creates a protected-mode program's INT 21h, custom something, and real-mode DOS code, kernel, and I/O, and from a real-mode DOS's real-mode request back to something that the program's code might not understand.

For a more complete discussion of DOS services and protected-mode services, see the second edition of *Extending DOS*, edited by Ray Duncan.

Just as Windows programs use DOS real-mode I/O, the real-mode DOS code can use DOS to stop around a limited, the system. But also Windows programs have transparent access to the standard documented DOS programs from the I/O and the kernel, and access to undocumented DOS, generally, not from real-mode. For example, status in protected-mode Windows programs can be done with an INT 21h AH=0x03, or more, regard to the real-mode DOS, which will only be located in a extended memory of some, or straight-forward for a Windows program to call an undocumented function on such as INT 21h AH=0x20 and the access a data structure, such as the CDS. The

Standard I/O Device I/O, DOS extenders built into Windows 3.0 and 3.1 do not support many undocumented DOS functions and data structures.

In addition to being a DOS extender that is a provider of protected mode I/O 21h services, Windows 3.x provides an implementation of version 3.09B of the DOS Protected Mode Interface (DPMI). DPMI includes services for generating real mode interrupts from protected mode, mapping real mode memory to a protected mode program's address space, and so on. As this chapter details a Windows programmer can use these services to access undocumented DOS because Windows programs can execute in protected mode, especially via DPMI. DPMI clients using DPMI consist of the IAS, DRV, and a portion of EMM386, for example, as a Windows program. Yet another version of IAS (IDRV) is available on CD-ROM 2, for example, as a Windows program. Yet another version of IAS (IDRV) is available on CD-ROM 2, for example, as a Windows program. Yet another version of IAS (IDRV) is available on CD-ROM 2, for example, as a Windows program.

Windows programmers using DPMI as a Windows extension IAS/IDRV could instead use some other extender. The more Windows API functions to access the undocumented DOS data structures, the more DPMI use has a more undocumented status than Windows. Windows itself uses a real mode DPMI interface (DPMI functions). There are vast numbers of undocumented Windows API calls. However, these are not addressed here since an entirely separate book, *Undocumented Windows*, deals with this already. So here, to place it in its undocumented aspect, Windows. What this chapter covers, however, is a different topic—Windows that in some way permit or advance undocumented DOS.

Chapter 1 showed that Windows itself uses many undocumented DOS and presented code to generate some DOS (M)API or a device driver VxD that, therefore, now one would go about doing real mode DOS (M)API or VxD. This is important because VxDs are the future of DOS systems programming. When a device driver or other kernel mode code is writing a real mode Windows application that accesses real mode DOS, it is not taking a real mode kernel mode code. Windows Enhanced mode uses it as DOS extenders and DPMI services, and of this chapter is to present a simple VxD that provides a simple access to real mode DOS (M)API functions that Windows does not otherwise support.

Some examples and books that have appeared since the release of Windows 3.0 have already discussed the use of DPMI to access undocumented DOS functions and data structures. However, Enhanced mode Windows is a different situation, namely, in addition to providing protected mode Enhanced mode programs, it also provides real mode for each DOS boot and several mode programs in a real 8086 mode, without using extended mode. Each DOS boot and program initialization DOS, including separate instances or copies of DOS, is a real mode, such as the Current Directory Structure and Swappable Data Area.

To examine the use of VxDs to capture a windowed real mode DOS, save in Windows, one visits at C:\BIN\BIN\WIN386\SYSTEM\FILE.MANAGER, sitting at C:\WIN31. In other words, Figure 1-1 shows a simple C++ client of the techniques this chapter presents is the combination of a real mode VxD and a stack of DOS. It is a real mode DOS program or even from a DOS program running under Windows.

It is a real mode subject, there is no sense to a mode and protected mode. This sounds "impaired" but upon the real mode client, it is not possible to use your program depends on will have already been ported to protected mode, so one could today many programs depend on other components which could be ported to protected mode. For the foreseeable future, the ability to get to real mode real mode code should be a necessary and invaluable skill. This is not especially so in Windows, other than VxD, in general, it is a necessary real mode engineering skill.

This chapter discusses the other topics, including writing simple Windows programs without VxDs, real mode windows of handle messages using DPMI to write protected mode DOS programs, real mode using I/O and services, real mode under Windows using I/O 21h to call services provided by Windows VxDs, writing Windows VxDs, and examining some Windows implements its DOS extender and DPMI server.

Before we begin a brief appeal to those DOS programmers who "don't do Windows" please don't expect a paper. We will be digging Windows down to some level of DOS systems programming, but it is not possible to do it. It is most common and part of DOS and opposite of Windows. In this chapter, VxDs, however, view Windows as little more than a protected mode version of DOS.

## Calling Undocumented DOS from Windows

In Chapter 2, we dealt about the intricacies of the LASTDRIV entry, but each time we ended up with a real-mode DOS program. I am sure the department responsible for this somewhat "exotic" wants us to turn LASTDRIV into a Windows program. This implies that we must deal with the usual how-getting-to-the-DOS-internal-data-structure-thing-protection-mode-How-does-it-proceed

## Windows and Printf()

**The first problem is what** to do with LASTDRIV's call to printf(). LASTDRIV only needs to produce a single line of output such as "LASTDRIV=0x1". The amount of output is exactly the same as that for the classic "Hello world" program, so writing the code shouldn't be too hard. However, all the standard books on Windows programming start by teaching you how to display "hello world" in Windows. The source code for this supposedly introductory program invariably takes about 80 lines of code and requires several other source files as well, just displaying these seven characters in Windows apparently requires calls to RegisterClass(), CreateWindow(), GetMessage(), DispatchMessage(), TextOut(), and so on, as well as an understanding of the concepts behind each of these key Windows API functions. It seems that, by agreeing to make the request to port LASTDRIV to Windows, we have bought into a big mess.

To survive as a PC programmer in the 1990s, you really do need to learn those Windows functions. However, that's a radically different place to start. The biggest initial stumbling block for DOS programmers moving to Windows is the misconception that, to write a Windows program, even the simplest Windows program, you must first learn how to register window classes, create windows, handle messages, and so on.

Well, it just isn't necessary. Anthony's *has Borland* with its *EasyWin* library and Microsoft with *QuickWin* made it easy to port simple DOS programs to Windows, but even more important, there are Windows API functions which are designed groups that you can use them to do an entirely self-contained, simple Windows application. Hello world? Despite what all the Windows programming books say, it's as simple as calling *MessageBox()* well, that and using Windows' weird *WinMain()* entry point instead of the *main()* entry point that's used everywhere else in the world.

```

/* gcc -MS hello.c */
#include "windows.h"

#pragma argsused
int PASCAL WinMain(HANDLE hInstance, HANDLE hPrevInstance,
                  LPSTR lpszCmdLine, int nCmdShow)
{
    MessageBox(0, "hello world", "My First Program", MB_OK);
    return 0;
}

```

That's it. Furthermore, the Windows *MessageBox()* function can produce an entire screen of output—just pass carriage returns in the string passed as the second parameter to *MessageBox()*. This makes it handy for tiny utilities like LASTDRIV.

Alternatively, you can take advantage of Windows' multitasking and its interprocess messaging to use another program, such as Notepad, as your "display engine." For example, a Windows program can use the *WM\_SETTEXT* message to blast text into some other window. This message can be sent to the other window with *SendMessage()*. For example, the following program launches Notepad and then sends the string "hello world" to Notepad's window. There's no intermed any text file.

```

/* bec -WS hello2.c */
#include "windows.h"

#pragma argsused
int PASCAL WinMain(HANDLE hInstance, HANDLE hPrevInstance,
LPSTR lpszCmdLine, int nCmdShow)
{
    if (WinExec(notepad.exe, SW_SHOWNORMAL) < 32)
        return FALSE;
    yield();
    SendMessage(GetFocus(), WM_SETTEXT, 0,
        (DWORD) (char far *) "hello world");
}

```

Using such primitive capabilities, it is easy to build Windows specific versions of functions such as printf(). There is nothing magical about printf() or any other stdio function. As EasyWin and QuickWin show it's possible to build stdio libraries on top of the Windows API. So why isn't there such a facility in Windows to begin with?

PRINTF.C and PRINT.H in Listings 3.1 and 3.2 provide a rather hokey (one tech reviewer says "a program that is a program" but I don't know how to do simple Windows utilities like the ones we build this way). A program that wants to use this facility should include "printf.h" and start off with `WM_OPEN_DISPLAY` to open display functions. The program can then track `open_display()`. In `PRINTF.C` just simply accumulate text until a call to show displays. To build a program, just work with the `RUNTIME` and `open_display` for Windows. For example, in Borland C++ use `bec.Whello.C` for the `Microsoft C++` was well as need a Windows `DEF` file automatically.

The `open_display` function normally calls `MessageBox`. But if there is more text than `MessageBox` can handle `PRINTF.C` uses `WinExec` to fire up a copy of Notepad. `PRINTF.C` then sends a `WM_OPEN_DISPLAY` message and the `SendMessage` function mentioned above. The program does not create an intermediary file; it shoots the text directly to memory to Notepad. For our program can send text to another shows incidentally that Windows is generally message based. `Message` is a low form of interprocess communication, not just an elaborate way of describing a function call.

### Listing 3.1. PRINTF.C

```

/*
PRINTF.C - simple output for small windows programs,
using MessageBox() or WinExec()/SendMessage()
from undocumented DOS, 2nd edition (Addison-Wesley, 1993,
*/
#include <stdlib.h>
#include <stdarg.h>
#include <string.h>
#include <windows.h>
#include <stdio.h> // for vsprintf
#include "printf.h"

#define BUF_SIZE 2048

static char *str, *app,
static unsigned cap, len,
static int lines;

BOOL open_display(char *appname)
{
    app = appname,
    cap = 128;

```

```

if (! (str = (char *) malloc(cap)))
    return FALSE;
*str = {lines = len = 0;
return TRUE;
}
/* Maximum number of lines that MessageBox will hold */
static int max_lines(void)
{
    TEXTMETRIC tm;
    HWND hWnd = GetActiveWindow();
    HDC hDC = GetWindowDC(hWnd);
    if (hDC == NULL)
        return 0;
    GetTextMetrics(hDC, &tm);
    ReleaseDC(hWnd, hDC);
    return (GetSystemMetrics(SM_CYFULLSCREEN) /
            (tm.tmHeight + tm.tmExternalLeading)) - 5;
}
BOOL show_display(void)
{
    if (lines <= max_lines())
        MessageBox(NULL, str, app, MB_OK);
    else
        notepad(str);
    free(str);
    return TRUE;
}
static BOOL append(char *s2)
{
    char *s3;
    if ((len == strlen(s2)) < cap) && strcat(str, s2))
        return TRUE;
    cap = len + 128;
    if (! (s3 = (char *) malloc(cap)))
        return FALSE;
    strcpy(s3, str);
    strcat(s3, s2);
    free(str);
    str = s3;
    return TRUE;
}
int nlines(char *s2)
{
    int c, n = 0;
    while ((c = *s2++) != 0)
        if (c == '\n')
            n++;
    return n;
}
int printf(const char *fmt, . . .)
{
    static char s2[BUF_SIZE];
    int len;
    va_list marker;
    va_start(marker, fmt);
    // Following uses vsprintf() rather than KERNEL wvsprintf()
    // because wvsprintf() requires that all %s par meters
    // be FAR; e.g., printf("x\n"), (char far *) "hi")
    len = vsprintf(s2, fmt, marker);
    lines += nlines(s2);
    va_end(marker);
    append(s2);
    return len;
}

```

```

)
BOOL notepad(char far *s)
{
    HWND notepad;
    HWND edit_ctrl;
    if (WinExec(notepad.exe", SW_SHOWNORMAL) < 32)
        return FALSE;
    notepad = FindWindow(NULL, "Notepad - (untitled)");
    edit_ctrl = GetFocus();
    SendMessage(notepad, WM_SETTEXT, 0, (DWORD) (char far *) app);
    SendMessage(edit_ctrl, WM_SETTEXT, 0, (DWORD) (char far *) s);
    return TRUE;
}
#endif TESTING
nt PASCAL WinMain(HANDLE hInstance, HANDLE hPrevInstance,
LPSTR lpszCmdLine, int nCmdShow)
{
    int i;
    open_display(GetSystemMetrics");
    for (i=0; i<37; i++)
    {
        printf("d\\dir\\n", i, GetSystemMetrics(i),
            i);
    }
    show_display();
}
#endif

```

**Listing 3-2: PRINTF.H**

```

/* PRINTF.H */
BOOL open_display(char *appname);
BOOL show_display(void);
#ifdef _cplusplus
extern "C"
#endif
int printf(const char *fmt, );
BOOL notepad(char far *s);

```

There are several alternatives to using this weird version of printf. As noted above, Borland C++ comes with a EasyWin library that provides C++ subroutines for Windows. To use EasyWin, you just take a DOS program—such that has main—instead of WinMain—as its entry point—and compile for Windows (for example, `bc -W hello.c`). To build a program with QuickWin in Microsoft C/C++, compile with the `/M` switch. In addition, the book *Undocumented Windows* comes with a more extensive set of libraries for Windows, WINLIB, which provides many features beyond those in EasyWin or QuickWin—including the ability to write event-driven programs while still using stdio functions.

There are also some versions of printf for Windows; it is easy to make LASTDRIV to Windows. For example, using TRNFILE as Listing 3.1, we hacked one of the versions of LASTDRIV.C from Chapter 2 so it started with WinMain—instead of main—Obviously, the program does the usual end-of-line (\r\n) set of calling. (N.B. The function `s2b` and indexing into `StrVars`. We might as well add a few more lines of output. The result is a Windows version of LASTDRIV.C, shown in Listing 3.3. Since so many programs here use the `getenv` function, it has been moved into a separate file, `ENVARS.C`, in Listing 3.3a.

**Listing 3-3: LASTDRIV.C for Windows**

```

/* LASTDRIV.C: gcc -W lastdriv.c printf.c */
#include <stdio.h>
#include <dos.h>

```

```

#include "windows.h"
#include "printf.h" // see listing 3-1, 3-2
void maybe_change_osmajmin(void)
{
    _asm mov ax, 3506h
    _asm xor bx, bx
    _asm int 21h
    _asm or bx, bx
    _asm jz osmajmin_ok
    _asm mov byte ptr _osmajor, bl
    _asm mov byte ptr _osminor, bh
osmajmin_ok:
}
/* get_sysvars() */
#include sysvars.c
#ifdef __cplusplus
extern "C" int PASCAL WinMain(HANDLE hInstance, HANDLE hPrevInstance,
    LPSTR lpszCmdLine, int nCmdShow);
#endif
int PASCAL WinMain(HANDLE hInstance, HANDLE hPrevInstance,
    LPSTR lpszCmdLine, int nCmdShow)
{
    BYTE far *sysvars;
    BYTE lastdrive;
    open_display("LASTDRIVE");
    sysvars = get_sysvars();
    printf("SysVars @ IFF\n", sysvars);
    maybe_change_osmajmin();
    printf("DOS E0 102uin, _osmajor, _osminor),
    lastdrive = sysvars[( _osmajor == 5 && _osminor == 0) ? 0x1B
        ( _osmajor == 2) ? 0x10
        /* otherwise */ 0x21);
    printf("LASTDRIVE=%2cin", 'A' + lastdrive - 1);
    show_display();
    return 0;
}

```

### Listing 3-3a: SYSVARS.C

```

/* SYSVARS.C */
BYTE far *get_sysvars(void)
{
    _asm xor bx, bx
    _asm mov es, bx
    _asm mov ah, 52h
    _asm int 21h
    _asm mov dx, es
    _asm mov ax, bx
    // return value in DX:AX
}

```

Compiling the program (for example, `lsc W lastdrv.c printf.c`) results in a Windows executable, `LASTDRIVE.EXE`. Given that the premise of this chapter is that Windows applications must do something special to access window-oriented DOS, and given that we have done nothing special (yet), you would fully expect `LASTDRIVE` to blow up or otherwise do something foolish when it runs under Windows.

Actually, though, `LASTDRIVE` works fine. Figure 3-1 shows that `LASTDRIVE` not only looks like a little Windows program, but the results look reasonable, too. In this configuration, `LASTDRIVE` is the correct answer.





```

printf("SDA @ %p\n", sda);
/* Following line gets wrong results! (Unless if UNDOSNGR 386,
   which supports 21/SD06 in protected mode, its installed) */
currdrive = sda[0x16];
printf("CURRDRIVE=%c (from SDA)\n", 'A' + currdrive);
doc_currdrive = get_doc_currdrive();
if (currdrive != doc_currdrive)
    printf("Something wrong! CURRDRIVE=%c\n",
          'A' + doc_currdrive);
}

#ifdef _cplusplus
extern "C" int PASCAL WinMain(HANDLE hInstance, HANDLE hPrevInstance,
    LPSTR lpzCmdLine, int nCmdShow);
#endif

int PASCAL WinMain(HANDLE hInstance, HANDLE hPrevInstance,
    LPSTR lpzCmdLine, int nCmdShow)
{
    open_display( CURRDRIVE );
    do_currdrive();
    show_display();
    return 0;
}

```

#### Listing 3-4a: GETSDA.C

```

BYTE far *get_sda(void)
{
    _asm push ds
    _asm push si
    _asm mov ax, 5d06h
    _asm int 21h
    _asm jc error
    _asm mov dx, ds
    _asm mov ax, si
    _asm jmp short done
error:
    _asm xor ax, ax
    _asm xor dx, dx
done:
    _asm pop si
    _asm pop ds
    // return value in DX:AX
}

```

CURRDRIVE really does seem just like FAS1DRIVE. Each program calls an undocumented function, and each program extracts a byte value from a DOS internal data structure whose address the function returns. While FAS1DRIVE worked though, CURRDRIVE produces totally wrong results, as Figure 3-2 shows.

Figure 3-2: CURRDRIVE Messes Up



Not only are the results wrong, CURRDRIV showed drive A as the current drive, whereas the disk-oriented function IOCTL returned C. But running multiple copies of CURRDRIV produces even stranger results. The address displayed for the SDA is different each time. For example, one run of the CURRDRIV showed an address of 11970320 for the SDA, whereas Figure 3-2 shows an address of 13870320 for the SDA, the segment for which looks wrong, anyhow. The SDA does not move around. Something is clearly wrong with the idea that Windows programs can just call any old undocumented Win32S function and access any old internal DOS data structure.

At the end of this chapter, see Listing 3-31, we'll write a VxD that changes Windows in such a way that CURRDRIV gets correct results. Without this VxD, however, CURRDRIV gets bogus results.

### The Dreaded GP Fault

\*But I don't ever call INE2H function 5106b. I've never used the SDA. The only undocumented DOS function I could see using is function 52b, and the only structure I'm interested in is SysVars. LISTDRIV.C showed that those work. End of story, right?

No, it's not so easy grabbing something other than LISTDRIV out of SysVars. A good example of the problem is the System File Table chain at offset 4 in SysVars. As Chapter 1 mentions, programmers WIN32M and related try to get the value of FILEN, reliably usually do so by walking the SE1 chain. SE1WALK.C (Listing 3-5) is a real mode DOS program that uses this technique.

### Listing 3-5: SFTWALK.C for Real Mode DOS

```
/* SFTWALK.C -- count FILES* by walking SFTs */

#include <stdio.h>
#include <stdlib.h>
#include <dos.h>

typedef unsigned char BYTE;
typedef unsigned short WORD;

/* get_sysvars() -- see Listing 3-3a */
#include "sysvars.c"

typedef struct _sft {
    struct _sft *next;
    WORD num;
    // other stuff not used here
} SFT;

main()
{
    BYTE far *sysvars = get_sysvars();
    SFT far *sft = *(SFT far *) sysvars[4];
    int files = 0;
    while (FP_OFF(sft) != 0xFFFF)
    {
        files += sft->num;
        printf("SFT @ %p -- %d files\n", sft, sft->num);
        sft = sft->next;
    }
    printf("FILES=%d\n", files);
}
```

In addition to printing out the FILES\* value, SFTWALK also displays the address and number of files for each SFT in the chain. For example:

```
SFT @ 0110:00CC -- 5 files
SFT @ 05EB:0000 -- 40 files
FILES=45
```

This output was correct. In this configuration, CONFER.SYS had FILES 45. However, before we've even turned this into a Windows program, just running the same `fs:dir` mode SFTWALK.EXE in a DOS box under Windows produced interesting results, claiming FILES 85 and showing an extra SFT with 10 entries:

```
SFT @ 0116:00CC -- 5 files
SFT @ 05EB:0000 -- 40 files
SFT @ 1295:0004 -- 10 files
FILES=85
```

We will explore this issue later. For now, the curious reader should check out the descriptions of the Windows SYSTEM INFORMATION files, starting in categories such as the *Microsoft Windows for Workgroups Resource Kit*.

To start porting SFTWALK.C to Windows, get started with Visual Basic. Visual Basic requires no change to the source code, merely recompiling with the Visual Basic Compiler. Or you can modify SFTWALK.C to use `PRINTF` from Listing 3-1; it makes no difference.

After `vbw`, when the Windows version of SFTWALK runs, it doesn't work as we normally see in LANSDRV.dll. Nor does it even produce incorrect results; the way LANSDRV.dll "casts" SFTWALK's productivity—general protection condition or GPF error—is Figure 3-3 shows. In Windows 3.0, the message refers instead to an "Unsuspected Application Error" (the default U.A.M.).

### Figure 3-3: Windows Version of SFTWALK GPF Faults

Application Error  
SFTWALK caused a General Protection Fault in  
module SFTWALK.EXE at 0001:01A4.

If SFTWALK is compiled with debugging symbols and run under a debugger such as Turbo Debugger or Soft ICE (Windows WINICE), the debugger catches the GPF fault and highlights the offending code in SFTWALK.C:

```
files = sft->num,
les bx, [ebp-8]
add cx, es:[bx+4] ; GPF fault occurs here!
```

The GPF fault is just an interrupt. INTEL didn't let the microprocessor send to an application what it has violated one of the many "rules of the road" in protected mode. These rules include write-processor mode protection. For instance, segments in protected mode are specified sizes, and if a program reads or writes to more than just that size, the processor decides that the program has violated protection and issues a GPF fault.

One rule of protected mode's strict compliance is that any address must be a multiple of segment register. It must be a specific number, called a *segment selector*, which corresponds to a processor's address range called a *descriptor*. Most faults are generated in contexts if used as `www.ssdw.com` pointers and for indexes. That is, the `www.ssdw.com` loaded into segment register such as `ES`, with `GP` and `INDEX` registers, which are most `www.ssdw.com` register sizes. In these spots, the address space is sparse. This is largely where the protection comes from—sparse protection is sparse, probably. Random draw might strike from a hat, and one of them is likely to be a valid selection. In practice, though, this protection works pretty well. In contrast to the havoc it would wreak in real mode, the following `www.ssdw.com` program, `x99-44` (100%), likes to be terminated with a GPF fault in protected mode:

```
main()
{
    unsigned char *fp;
    for ( ; ; )
    { // see Listing 2-10 for MK_FP()
        unsigned far *fp = MK_FP(rand(), rand());
```

```
*fp = 606;
```

The error is that protected mode pointers must correspond to descriptors. These descriptors are kept in segments—maintained generally by an operating system such as Windows—on a base-by-the-top basis with each segment being headed by a segment register. A descriptor contains information about a segment, such as the specific size of the segment. The size actually is stored as the *limit* (the segment's maximum offset), that is, as the last valid byte offset within the segment. Besides the segment's size, a descriptor also contains a segment's attributes: code, data, read, write, and so forth, and its *base address*.

This notion of base address is crucial for this chapter. We will rely on it heavily to get to unaligned memory (DOS data structures from protected mode). Recall that a real mode address  $xxxxxxx$  corresponds to the memory location  $xxxx * 10 + xxxxx$ , so that  $10070000$ , for example, points to the  $xxxx = 10070$  protected mode address. A small but crucial change: There are still  $xxxxxxx$  pointers, not the  $xxxx$  part means something entirely different.

It's a nice protected mode address, a collection of instructions from memory addressing. A selector, for example, is  $10070$ , for example, is a key into a descriptor table. This value is not in any way related to physical memory location  $10070h$ . Actual memory addresses are, instead, found inside the descriptor, in the address field, and are actually stored to get to a memory location such as  $40070h$ . A protected mode pointer such as  $10070000$  would do you absolutely no good at all. Instead, a normal real mode selector whose descriptor had a *base address* of  $10070h$ . The value of the selector would be an index into the selector table, which represents an index into a descriptor table.

But a pointer to a table that an address is not accessible by physical memory address, that is, a value that a processor cannot use. If paging is used on the 80386, higher, it is instead *virtual address*. This corresponds to real mode selector tables, and is another level of indirection. Page tables convert the protected mode address to a indication that a page is present, this is the only virtual memory type. Exploring Normal/Paged Virtual Memory in Windows Enhanced Mode. *Microsoft Windows Journal*, February, 1992.

Fortunately, OS/2 makes this transparent. Programmers can access the segment registers just like in real mode, and the processor takes care of converting out to the correct part of memory. It is also fortunate that the hardware has a series of capabilities that the processor contains several caches, a cache array, and fast access to real mode tables, so you can do some virtual-to-real memory.

Now, some real mode programs (ASNDRV, C:\RRDRV, and SFWALK) do not or test for real mode, yet will get DOS, will still run, results. ASNDRV works properly. SFWALK, C:\ASNDRV, and C:\RRDRV, does not. What, but instead produces bogus results. What can we conclude about Windows 95's extended support for real mode (EEM)?

- ASNDRV works, so the  $SwVars$  pointer must be valid. The program didn't GP fault, so  $SwVars$  exists, and the real mode pointer,  $INT_21h$   $AF$   $82h$  returned this pointer, so in Windows protected mode, this must be a real protected mode pointer to  $SwVars$ .
- C:\RRDRV, does not GP fault, so the  $SDV$  pointer must at least be legal. It's a legitimate protected mode pointer, and the results are bogus. Conclusion: it must actually point, not to the  $SDV$ , but to some real mode. Thus,  $INT_21h$   $AF$   $5000h$  in Windows protected mode returns protected mode pointer, which can be  $SwVars$ .
- SFWALK GP faults because the  $SDV$  pointer is invalid, from protected mode. That means it must be a real protected mode pointer. But, it's a real mode pointer out of  $SwVars$ , to which  $INT_21h$   $AF$   $82h$  returns a real protected mode pointer. Could point in Windows protected mode,  $INT_21h$   $AF$   $82h$  returns a protected mode pointer to  $SwVars$ , but any pointers inside  $SwVars$  use real mode, and will need to be translated.

The implications of the earlier discussion on descriptors and selectors (we get a `cs` and `index` membered DOS data structures from protected-mode Windows should now be clear. To get at the SEI, for instance, you can't load the real mode address of the SEI into ESI/EAX and expect it to work. Instead, you must allocate a descriptor for the real mode address of the SEI and use the descriptor's base address field, and increase whatever selector corresponds to that descriptor as the `ss` portion of your protected mode `ss:0000` pointer. We're getting a good bit ahead of ourselves here—within the Windows API the function `AdLocSelector` allocates both a descriptor and a selector, the formerly undocumented function `SetSelectorBase` sets the base address. We will be relying heavily on these functions in this chapter.

### A DPMI Shell

To gain a better understanding of what's going on and what we can do about it, let's forget a world of Windows programs for a moment and look at DPMI programming. Using the DPMI services that Windows Enhanced mode provides, real-mode DOS programs can switch into protected mode. While Windows Standard mode also provides DPMI, it is only a *residual* Windows program, not to programs running in the Standard mode DOS box. Memory managers such as 386MAX and `svga` add in, QEMM, also provide DPMI services.

Before we launch into using DPMI, it's worth citing an error in the Microsoft Windows SDK documentation. The *win* brief SDK Assessment of DPMI Windows Applications with MS-DOS Functions *Programming Reference Manual* (Chapter 26) states that Windows 3.0 and later support DPMI LIB. This is not true. Both Windows 3.0 and 3.1 support version 0.9 of the DPMI specification. Further, not all DPMI functions are supported, even those of Standard mode. On the other hand, Enhanced mode does support a few DPMI functions—commonly available VxDs. The SDK also claims that all seven DPMI functions are reserved for Windows applications. This certainly is true, but only when the new 3.11 functions are included; they are stable across all the different DPMI implementations.

We can use DPMI to take DOS versions 4.1, ANSDRIVE, CURDRIVE, and SHIMALK and execute them in protected mode, so running in protected mode and stacking a character study DOS, we can separate the Windows-specific issues from the more general protected mode issues. Comparing the behavior of the protected mode DOS versions of these programs with the behavior of the Windows versions also yields some support for the above Windows.

A minimal real mode DOS program can easily switch itself into protected mode using only DPMI call INI 21h:AN 1687h. Moving the DPMI-specific code into a separate `shell` which makes the `call INI 21h:AN 1687h`. Moving the DPMI-specific code into a separate `shell` which makes the `call` and then starts off executing the program, can make using DPMI easier. In DPMISHELL, Listing 3-6, `main` calls the function `runmain`, then uses DPMI call 21h:AN 1687h to protect real mode, and then calls `main`, which is DPMISHELL's application such as ANSDRIVE to supply real mode and `main` (real mode) and to `runmain` DPMISHELL (see Listing 3-7). Besides DPMISHELL, the program also needs to deal with `Ctrl-C`, a run-time language module which assists with `Ctrl-C` handling, see Listing 3-8.

### Listing 3-6: DPMISHELL

```

/*
DPMISHELL.C
Shell to run a simple C program in protected mode under DPMI
Andrew Schulman, February 1993
from Undocumented DOS, 2nd edition (Addison-Wesley, 1993)
(Changed version of DPMISHELL from MSJ, Oct. 1992, Dec 1992)

Boss Ctrl-C handling; see CTRL_C.ASM
Must be compiled with a model to use C run-time library
Calls real_main(), switches into protected mode, then calls pmode_main()

To build a DPMI app:

```

```

bcc -2 -DDPMI_APP foo.c dpmish.c ctrl_c.asm
*/
#include <stdlib.h>
#include <stdio.h>
#include <dos.h>
#include "dpmish.h"

void _dos_exit(int retval)
{
    _asm mov ah, D1ch
    _asm mov al, byte ptr retval
    _asm int 21h
}

/* Call the DPMI Mode Detection function (INT 2fh AX=1686h) to see
// if we are "a.ready" running in protected mode under DPMI.
// See 3.1 BDK Device Driver Adaptation Guide, pp. 585-586
int dpmi_present(void)
{
    unsigned _ax,
    _asm mov ax, 1686h
    _asm int 2fh
    _asm mov ax, ax /* 2f/1686 ret 0 if DPMI is present */
    return (!_ax), /* turn it around so 0 if "not" present */
}

// Call the DPMI function for Obtaining the Real to Protected Mode
// Switch Entry Point (INT 2fh AX=1687h), to determine if DPMI is
// available and, if so, switch into protected mode by calling
// the Switch Entry Point. See DPMI 0.9 spec.
int dpmi_init(void)
{
    void (far *dpmi)(),
    unsigned hostdata_seg, hostdata_para, dpmi_flags,
    _asm push si
    _asm push di
    _asm mov ax, 1687h /* test for DPMI presence */
    _asm int 2fh
    _asm and ax, ax
    _asm jz gotdpmi /* if (AX == 0) DPMI is present */
    _asm jmp nodpmi
gotdpmi:
    _asm mov dpmi_flags, bx
    _asm mov hostdata_para, si /* paras for DPMI host private data */
    _asm mov word ptr dpmi, di
    _asm mov word ptr dpmi+2, es /* DPMI protected mode switch entry point */
    _asm pop di
    _asm pop si
    if ( dos_al ocmeemhostdata para, &hostdata_seg) != 0)
        fail("can't allocate memory");

    dpmi_flags &= 1, /* this is a 16-bit protected mode program */

    /* enter protected mode */
    _asm mov ax, hostdata_seg
    _asm mov es, ax
    _asm mov ax, dpmi_flags
    (far *dpmi)(),
    _asm jc nodpmi /* carry set if error */
    /* in protected mode now segment registers changed */
    return dpmi_present(); /* double check */
nodpmi:
    return 0,
}

void dpmi_setprotect(int intrno, void (interrupt far *func)(void))

```

```

    _asm mov ax, 0205h
    _asm mov bl, byte ptr intno
    _asm mov cx, cx /* word ptr func+2 */
    _asm mov dx, word ptr func
    _asm int 31h
}

// INT 23h handler under DPMSI can't do the usual DOS INT 23h stuff
// needs to be on page locked with 31:0600**
// problem: compiler has hard-wired (real mode) loadsd!
#if 1
// pull in CTRL_C.ASM (Listing 3-8)
#ifdef __cplusplus
extern "C" void interrupt far ctrl_c(void),
#else
extern void interrupt far ctrl_c(void),
#endif
#endif
int ctrl_c_hit = 0;
void interrupt far ctrl_c(void)
{
    ctrl_c_hit++;
}
#endif
main(int argc, char *argv[])
{
    int ret,
    // actually, if already in pmode, real main() still runs
    // (this is for debugging under Z80/DOS-Extender)
    // if (real_main(argc, argv) != 0)
    return 1,
    fflush(), // flush all buffers before switch into protected mode
    // so I/O redirection works properly
    if (dpmi_present())
        printf("Already in protected mode\n");
    else if (dpmi_init())
        printf("Switched into protected mode via DPMSI\n"),
    else
        fail("This program requires DPMSI"),
    // now in protected mode, segment registers have changed
    dpmi_setprotect(0x23, ctrl_c), // install Ctrl-C handler
    ret = pmode_main(argc, argv), // call the application's pmode_main
    fflush(), // flush all buffers before exiting
    dos_exit(ret), // must exit via 21 4C!
}

```

### Listing 3-7: DPMSH.H

```

/*
DPMSH.H -- see DPMSH.C
Andrew Schulman, February 1993
from "Undocumented DOS", 2nd edition (Addison-Wesley, 1993)
*/
#ifdef __BORLANDC__
#ifdef __SMALL__
#error DPMSH requires small model
#endif
#else /* Microsoft C */
#ifdef M_I86M
#error DPMSH requires small model

```





**Listing 3-9: DPMITEST.C**

```

/*
DPMITEST.C
bec -DPMI_APP -Z dpmitest.c dpmish.c ctrl_c.asm

Sample output
CS=165Ch DS=1795h
Switched into protected mode via DPMI
CS=0097h DS=00Bfh
*/
#include <stdlib.h>
#include <stdio.h>
#include <dos.h>
#include "dpmish.h"

void print_regs(void)
{
    struct SREGS s;
    sregread(&s);
    printf("CS=%04Xh DS=%04Xh\n", s.cs, s.ds);
}

void fail(const char *s, ...) { puts(s), _dos_exit(1); }
int real_main(int argc, char *argv[]) { print_regs(), return 0; }
int main(int argc, char *argv[]) { print_regs(), return 0; }

```

Naturally, the new protected-task selectors that DPMI creates and loads into CS, DS, and SS on the program's behalf have base addresses that correspond to the program's virtual real-mode segment addresses. In the output from DPMI, I did show at the top of Listing 3-9, selecting DPMI's trap opened to have a base address of 16540h and selector 00Bfh (real-mode address 0:1650h). You can verify this in `WINDBG` or in the command-line message queue with the DPMI user Segment Base Address function (`INT 31h`). As in the case with DPMISH, only a program's protected-mode base address on the program is naming its protected-mode. The environment in which you execute a program is responsible for that of a Windows program, except that the program is in character mode and the various versions of third-party such as `printf` still work. This is why DPMISH is handy to give the system a set of libraries such as `EasyWin`, `QuickWin`, and `WIN32` in the case of combining together a Windows `printf()` from the `MessageBrot()` function.

But the DPMISH program only runs in library output as subtle as it is complete: a small model. Many of these old programs end up making `INT 21h` calls thereby using Windows' protected-mode DOS extension. DPMISH simply assumes that the presence of DPMI also means a user protected-mode `INT 21h` is present. This questionable assumption is also well known: a practice that works. But any small model program will never see the program or use DPMISH's going to have its segment registers changed, not to mention that the program cannot bring onto any of pointers after executing the `fdopen` call to switch into protected mode. The `printf` library undoubtedly does have some pointers for buffers and so on, so use a small model to make sure these won't be far pointers with segments or frames that will become invalid after the jump into hyperspace (oops, protected mode).

**Trying Out Undocumented DOS from DPMISH Programs**

We now take the Windows version of `LASTDRIV` and modify it to also run under DPMISH. We take the code that `LASTDRIV` executes inside of `WinMain` or `main` and move the code to a separate function, `do_lastdrv`. The DPMI APP version of `LASTDRIV` then calls `do_lastdrv` to open both real mode and protected mode. Listing 3-10 shows this new version of `LASTDRIV`.

**Listing 3-10: LASTDRIV for DPMI or Windows**

```

/*
LASTDRIV.C

```

```

Windows: bcc -M lastdriv.c printf.c
DPRI.   bcc -DDPRI_APP -2 lastdriv.c dpnish.c ctrl_c.asm
*/

#include <stdlib.h>
#include <dos.h>
#ifdef DPRI_APP
#include <stdio.h>
#include "dpnish.h"
typedef unsigned char BYTE;
#else
#include "windows.h"
#include "printf.h"
#endif

void maybe_change_osmajorin(void)
{
    _asm mov ax, 3306h
    _asm xor bx, bx
    _asm int 21h
    _asm or  bx, bx
    _asm jr  osmajorin_ok
    _asm mov byte ptr _osmajor, bl
    _asm mov byte ptr _osminor, bh
osmajorin_ok:
}

/* get_sysvars() -- see Listing 3-3a */
#include "sysvars.c"
void do_lastdrive(void)
{
    BYTE far *sysvars;
    BYTE lastdrive;
    sysvars = get_sysvars(),
    printf("SysVars @ 1fp\n", sysvars),
    maybe_change_osmajorin();
    printf("DOS @ 102u\n", _osmajor, _osminor),
    lastdrive = sysvars[( _osmajor == 3 && _osminor == 0) ? 0x18 :
        ( _osmajor == 2) ? 0x10 :
        /* otherwise */ 0x21];
    printf("LASTDRIVE=Zc\n", 'A' + lastdrive - 1);
}

#ifdef DPRI_APP
void fail(const char *s, ...) { puts(s), _dos_exit(1), }
int real_main(int argc, char *argv[])
{
    printf("In real mode:\n");
    do_lastdrive();
    printf("\n");
    return 0;
}

int pmode_main(int argc, char *argv[])
{
    printf("In protected mode:\n"),
    do_lastdrive();
    return 0;
}
#else
// Windows program
#ifdef _cplusplus
extern "C" int PASCAL WinMain(HANDLE hInstance, HANDLE hPrevInstance,
    LPSTR lpszCmdLine, int nCmdShow),
#endif
#endif
int PASCAL WinMain(HANDLE hInstance, HANDLE hPrevInstance,

```

```

LPSTR (pszCmdLine, int nCmdShow)
{
    open_display("LASTDRIVE"),
    do_lastdrive(),
    show_display(),
    return 0;
}
#endif

```

The DPMI version of LASTDRIVE works just as the Windows version did. There seems to be some commonality between Windows programs and protected mode DOS programs. But when the DPMI version of LASTDRIVE gets its same exact LASTDRIVE\_MANAGER command, that is from real mode—the address is INI.219, Function 32b returns for SYSVARs in protected mode is different. The program's output looks like this:

```

C:\JNDOC2\CHAP3>lastdrive
In real mode:
SysVars @ 0116 0026
DOS 6.00
LASTDRIVE=M

Switched into protected mode via DPMI
In protected mode:
SysVars @ 00AF.0026
DOS 6.00
LASTDRIVE=M

```

Note that SYSVARs was located at 0116:0026 in real mode but at 00AF:0026 in protected mode. (Did SYSVARs move? No, of course not. Instead, there is just another confirmation of what you saw before. INI.219, function 32, in Windows returns a pointer to protected mode properties for SYSVARs. The stringger such as WINCL1 is used to get LASTDRIVE. It causes that 00AF:0026 in fact is protected mode selector and the base address 0116:0026. In other words, it addresses exactly the same memory, virtual mode pointer 0116:0026. Below 11F1 is the WINCL1 command for inspecting selectors. 11F1 refers to the Local Descriptor Table.

```

r[di 00af
00af  Base16 Base=00001160 Lsm=0000ffff CPL=3 P #W

```

WALLYabc use DPMISH or the CURRDRIVE in protected mode DOS program. It's not about showing the code here, because I cannot see code. It lists 8 4 more the same way that Listing 3.10 relates to Listing 3.3. We just stick the code in WinMain, so that mode can be seen.

This protected mode DOS version of LASTDRIVE gets an incorrect answer, and this is Windows version 4.0. Oddly, error 3 is same as it appears even from the CURRDRIVE code. I called it DPMI server version 5.06MAY 6.0. The important lesson was, and this is a common thing to do, and DPMISH is that protected mode DOS programs running under Windows cannot test the character mode Windows programs. The environment for protected mode is a DPMISH program is a for the WinMain. A DPMISH program, or any other DOS program, that switches itself into protected mode with INT 31H (0x168) is a part of the character mode Windows program. Various, this amount of time point is, as you can see, that Windows programs are without any form protected mode DOS program. Windows is just a protected mode extension to DOS. Still wondering why we're talking so much about Windows in a book on DOS?

### The Windows DOS Extenders

Since MS-DOS is not a protected mode operating system, how can there ever be several layers of protected mode DOS programs? The answer is a protected mode version of DOS must be *added*.

The basic role of the DOS extender. In Windows Enhanced mode, the DOS extender is named in the DOSMIBs.VxD, which we discussed at length in Chapter 4, and which passes its

way use of user-oriented DOS. The DPMI server in Enhanced mode—that is, the code that serves functions such as INT 21h AX=400h and the DPMI INT 31h API—is provided by the Virtual Machine Manager (VMM) and DOSMGR, both part of WIN386.EXE/DOS386.EXE in Ch. 6. In Standard mode, DOS 5.1A contains both the DOS extender and the DPMI server.

The documentation for the DOS extenders in Windows is the scant four-page chapter on “Windows Applications with MS-DOS Extensions” in the Windows 3.1 SDK *Programmer's Reference, Volume 1*. Chapter 14 defines the largely inaccurate description of DPMI (noted earlier, the documentation also provides a too-short list of DPMI functions available to Windows programs). This chapter then lists both “fully supported” and “partially supported” DOS functions. The section on partially supported functions contains some important notes about DOS 11 calls, INT 21h AH=44h. There is no mention of any support of unsupported ROM BIOS functions. No mention is made one way or the other of important DOS services such as MS-DOS 4.0. The documentation does make the important point that “if a software application tries to completely register based on hardware pointers, segment registers, or stack parameters, this is not allowed even with Windows running in protected mode.”

Recalling the SDK documentation, there was one other Microsoft document, “MS-DOS API Extension DPMI History” (October 31, 1990), which devotes about thirty pages to this same subject. The 300+ chapter in the SDK appears to have been boiled down from this more extensive document. For example, the 1990 document discusses 32-bit DOS extenders. The DOS file read and write calls (INT 21h AH=31h and 40h) have the second register (CX) extended to 32 bits, allowing 32-bit programs to perform DOS file I/O on more than 64k bytes at a time. The Enhanced mode DOS extender provides the 32-bit support, one of our tech reviewers tells us this 32-bit support is “totally bogus.” The document also discusses the Ctrl-Break (INT 23h) and critical error (INT 24h) handlers with Windows. Finally, there is a somewhat discussion of ROM BIOS functionality, including low-level disk access (INT 13h) in protected mode. Microsoft should update this documentation for Windows 1.0 and make it available to all developers.

Given the paucity of documentation for the Windows DOS extenders, we set out on our own trying to figure out what they support. This could, of course, also take protected mode DOS. One good test, of course, what undocumented functions the DOS supports. From LANHDRV and CURDRIV, we've gathered that Windows supports functions 82h in protected mode and does not support Function 8300h. (How does one say “And what does ‘support’ and ‘not support’ mean here, anyway?”)

The next question is, “How can I find out?” The answer helps clarify this question: UNDOC can be built as a Windows program, using EasyWin, QuickWin, or WIN32, or hacked sagely to use PRINTC from Listing 3.1. The program acts as a proxy on the addresses of some important DOS internal data structures as returned from various INT 21h functions (as found within SysVars).

### Listing 3.11: UNDOC.C

```

/* UNDOC.C */
#include <stdio.h>
#include <dos.h>
#include <stdio.h>

/* get_sysvars() see listing 3.1a */
#include "sysvars.c"

void far *get_sft(void)
{
    unsigned char far * far *sysvars = get_sysvars();
    return *(sysvars+1); // sysvars[6]
}

void far *get_dpb(void)
{
    unsigned char far * far *sysvars = get_sysvars();
    return *sysvars; // sysvars[0]
}

```

```

void far *get_cds(void)
{
    unsigned char far *sysvars = (unsigned char far *) get_sysvars(),
    return *((void far * far *) &sysvars[0x16]);
}

/* get_sda() -- see Listing 3-4a */
#include "get_sda.c"

unsigned char far *get_indos(void)
{
    _asm mov ah, 34h
    _asm int 21h
    _asm mov dx, es
    _asm mov ax, bx
}

main()
{
    printf("SysVars @ %p\n", get_sysvars()); // supported!
    printf("CBS @ %p\n", get_cds());
    printf("SFT @ %p\n", get_sft());
    printf("DPB @ %p\n", get_dpb());
    printf("SDA @ %p\n", get_sda()); // not!
    printf("INDOS @ %p\n", get_indos()); // now doc!
    if (get_indos() != (get_sda() + 1))
        printf("Something funny with SDA and INDO5!\n");
    return 0;
}

```

Figure 3-4 shows output from UNDOC. Because the INDO5 flag is kept at offset 16 in the SYVARS area, the address that get\_sda() INI 21h AX 5400h ret instruction should be one less than the address of the INDO5 flag. It was finally documented in DOS 5.0 (see Chapter 4), so it is reasonable to expect that the Windows DOS extender would support this function, but UNDOC's output shows that the get\_sda() and get\_indos() return values are inconsistent. As we suspected, Windows doesn't support function 5400h.

**Figure 3-4: UNDOC Under Windows**

```

SysVars @ 1007 0026
CBS @ 0692 0000
SFT @ 0116 00CC
DPB @ 0116 1368
SDA @ 1E67 0320
INDOS @ 1007 0321
Something funny with SDA and INDO5!

```

That UNDOC is able to successfully extract values from SysVars indicates that a mixed reality get\_sysvars() INI 21h AX 5400h is not so returning valid and protected mode pointers. The pointers that UNDOC displays are, yes, valid, but they are not real mode pointers; the WIN32 debugger confirmed, for example, that the CBS was at 0692 0000. The level of support for function 521h is, if necessary, somewhat uneven.

Since, however, not whatever the level of support Windows provides, or the functions that UNDOC calls, the program does not fail. A GP fault occurs when a kernel-mode address is used as a pointer. UNDOC displays valid pointers which are useful in protected mode, but does not *display* them; there is a difference between printing "1p", "lp", and "fp".

It is instructive to run UNDOC in enhanced mode under the debug version of Windows. Recall from Chapter 1 that Windows Enhanced mode sets up WIN386 EXT, which is a collection of VxDs. The Enhanced mode DOS extender is part of the DOSMGR.VXD which is responsible for running in WIN386 EXT. The Windows Device Driver Kit contains a debug version of WIN386, which dis-

pass any error messages and which is newer than the retail version of WIN386. If you run UNDER under the fake WFN386 and in turn run Windows under a low-level debugger such as WIND or Microsoft's WDI386, the DOSMGR.DOS extender issues the following message:

```
WARNING: DOS INT 21 call AK-5006 will not be translated
Use of this call is not supported from Prot
mode applications
```

DOSMGR.DOS displays this message if you run a protected-mode version of the CURRDIR program.

Of course, if you use a Borland C++ compiler (including UNDER C++ or any other program of the fake products) this same message is emitted, but that Borland C++ assumes that all DPMSI services support INI 21h. Windows 3.11 runs in protected mode. It is not clear why Borland's compilers function in protected mode. Windows 3.11 is not running in the wrong place. Perhaps this is why Borland C++ compilers cannot be used to write Windows 3.11 software without DPMSI boot. As you've seen, actually, it is not the compiler's lack of support for protected mode that is the problem.

At the same time, C++ is making several bogus assumptions: that merely support for INI 21h in Windows 3.11 is protected mode. When running under DPMSI it assumes that all of INI 21h is supported. In protected mode, Windows 3.11 does not support INI 21h. In fact, DPMSI, if it supports INI 21h, does not support protected mode. INI 21h is not part of the DPMSI specification. If you use a Windows 3.11 real-mode DOS extender, it is somewhat likely that it is not a DPMSI service. As noted above, in Windows 3.11 real mode the DPMSI service is DPMSI and the DOS extender is in DOSMGR.DOS. In protected mode, under the fake WFN386, also produces many occurrences of the message "INT 21 call AK-5006 is not supported." So apparently BC++ is making INI 21h calls in protected mode too (why?).

In practice, it is a fairly reasonable assumption that the presence of DPMSI means that you also have protected mode INI 21h. As we saw, DPMSI also makes this assumption, which is why you can use it to find functions supported by the C runtime library functions that demand an INI 21h service. The message is strictly a false statement. For example, calling printf() in protected mode causes only the DOS Write File function, INI 21h AH 40h. In Windows 3.11, the DOS Write File function is not a DPMSI service. In fact, DPMSI and Borland C++ simply assume that the code function is not located on a function address beyond the limit of the assumption is valid only. And yes.

Of course, the assumption is not a good assumption. As we will see in more detail shortly, DPMSI provides services such as INI 21h AH 40h. Service Real Mode Interrupt to call down to real mode. Any program that uses DPMSI cannot be aware that it has real-mode support from INI 21h in protected mode. Can instead call INI 21h real mode. In fact, the first set of *Undocumented DOS* (see *UNDOS* command) DPMSI supports a DPMSI that is actually far more cautious than any other DPMSI. In fact, among the C runtime library's printf() and there are no other C runtime library's DOS extender in Windows 3.11/DPMSI included the functions printf(), printf(), printf(), and printf() which called down to real mode INI 21h using INI 21h AH 0300h.

Given the amount of Microsoft's documentation for the Windows DOS extenders, it is probably a good idea to be cautious and more cautious not only in using INI 21h functions that Microsoft documentation says are fully supported. This would mean that even though our code says that some function is supported, it would be safer not to trust this function. We will find the correct version of get services later in this chapter, see Listing 3.18. Certainly, get services must be changed (see Listing 3.27).

### Inside the DOSMGR.DOS Extender

It is how we can find out exactly what function one of the Windows DOS extenders supports. We have discussed how to use the *UNDOS* command. Now we will actually see what the Windows DOS extender is doing at the code, of course. In Chapter 4, we examined the DOSMGR.VXD to see

which undocumented DOS functions it calls. You can also see which protected-mode code (and protected DOS calls) it supports, and how. This will help determine what can be supported or not supported. For standard tasks you would need to examine DOSAPI (here, see `dosapi.asm`) in Enhanced-mode DOS extender, DOSMGR. As in Chapter 1, DOSMGR is also described using the author's Windows Source product from A Communications.

DOSMGR is a piece of 32-bit protected-mode code, running at Ring 0, which is being run at privilege level 0. Windows applications (making INE 21h calls) are the ones that are actually 16-bit and run at Ring 3, which is the most privileged level. This is not possible, so we know. How does DOSMGR arrange to make its protected-mode INE 21h handler to intercept Windows program opens a file, for example?

Among the many other things it does during the Sys Critical Init, etc., DOSMGR installs a protected-mode INE 21h handler, using the VMM Set VM and Vector function. The address DOSMGR passes to Set VM of Vector is that of a callback which prepares for a call to INE 21h by protected-mode code up using the various registers and code in the VxD. As shown in Listing 3-12, the VMM Allocates PM Call-back function creates the callback.

### Listing 3-12: Installing the Enhanced Mode DOS Extender

```
06f5c    mov     eax, 21h                ; INT 21h
06f71    mov     esi, PR_INT21         ; address of DOS extender (see Listing 3-13)
06f76    call   SetPRIntVect          ; PR = protected mode
;
; SetPRIntVect
06f88    mov     edx, esi              ; reference data = 21h
06f8b    VMRCall Allocate_PM_Call_Back
06f93    mov     ecx, esi              ; segment offset of callback
06f95    xchg   eax, edx              ; [EAX = interrupt number 21h]
06f96    movzx  edx, ds                ; EDX = handler offset
06f99    shr    ecx, 10h              ; CX = handler segment
06f9c    VMRCall Set_PM_Int_Vector
```

The protected-mode INE 21h handler, like DOSMGR itself, may well be created by a Windows program, via a DOS program running in protected mode, may be the `CPMISH` or `ztools` makes an INE 21h call. As shown and stated in Listing 3-13, the code in `SetPRIntVect` is very simple.

### Listing 3-13: The Enhanced Mode DOS Extender

```
PR_INT21
00c48    VMRCall SetPRIntVect
00c4e    movzx  eax, [ebp, Client_AH]   ; get INT 21h func #
00c52    cmp    eax, 0Ch                ; is it one we know?
00c55    ja     short UNKNOWN_FUNC
00c57    movzx  edx, byte ptr FUNC_TYPE[edx] ; get func type
00c5e    mov    edx, dword ptr XLAT_MACRO[edx*4] ; get VBMMGR script
00c65    VxDCall VBMMGR_xlat_API       ; run the script
00c68    UNKNOWN_FUNC
00c6b    mov    edx, INT21_DEFAULT     ; default script
00c70    VxDCall VBMMGR_xlat_API       ; run the script
```

In other words, the DOSMGR's protected-mode INE 21h handler knows (see it the function number specified in AH) is also that is, DOSMGR uses the AH function context is provided in the first of two bytes. This first byte groups the DOS functions into 16-bit sub-groups, and of course their input is a 16-bit register size. The function's signature is in whether or not they need any special treatment.

For example, `Find` and `30h Create Directory`, `31h Delete Directory`, and `31b Open File` all take strings in DS:DX, they return values in AX, and `31c` is the carry flag. DOSMGR has a few pro-

ected mode calls to all three functions in the same way. It takes the string from DS:DX, copies it to a transfer buffer in conventional memory, switches to real mode (actually V86 mode), puts the real mode address of the transfer buffer into DS:DX, issues the INT 21h call, and then jumps back to protected mode. Given that the many different INT 21h functions actually require only a handful of functions, real mode DOS functions seem to assemble them according to type. This is very similar to what one would expect to see in Real-Mode Procedure Call RPC over a network. In many ways, writing a DOS extender similar to RPM programming.

DOSMMGR also takes this INT 21h function type and uses it to index into a second table, which is known as the *translation table*. The V86MMGR.VxD (also part of WIN386.EXE) provides an API to a state service (documented in the Windows DDK) which helps build DOS extenders and other VxDs in a similar fashion. The V86MMGR service is actually a set of script macros defined in the .CHK files in the V86MMGR.INC and the V86MMGR.Mac API function is a LHA interpreter that runs these script macros. They simply do the job code over at Microsoft. V86MMGR also provides access to the state service. The state service, in conventional memory. The MMHMMINI (386F0h) VxD File Size setting (documented in the *Windows 3.0 Backup Resource Kit*) controls the size of this buffer; its default size is 8k bytes (4k in Windows 3.0).

DOSMMGR Uses V86MMGR.VxD API script macros to implement the DOS extender. The DOSMMGR.VxD also uses these macros, for example, to support INT 13h disk calls in protected mode. As I will explain, consider a function that requires special handling, such as INT 21h VxD 30x. The V86MMGR can just pass this call passed down ("reflected") to DOS. DOSMMGR does so using the following very simple V86MMGR macro:

```
Xlat_API_Exec_Int 21h
V86MMGR INC defines Xlat_API_Exec_Int like so:
```

```
Xlat_Exec_Int EQU 00h
Xlat_API_Exec_Int MACRO Int_Number
    db Xlat_Exec_Int
    db Int_Number
ENDM
```

Even the DOS extender's look for the default INT 21h handler is simply the two bytes 00h 21h, the address which DOSMMGR passes to V86MMGR.Mac API. These macros can make disassembly of DOSMMGR (and DOS) somewhat easier. A Common macro (has a VxD job) that translates these macros.

What does V86MMGR.Mac API do with these bytes? It uses the first one as an index into a table of 21h different macros, calls just the first one, leaving the specific handler to deal with any other jobs. It does this by the handler is "returning" out:

```

V86MMGR Xlat_API    proc near
00158    push    eax
0015c    movzx   eax,byte ptr [edx]    ; [EBX + pointer to script
0015f    inc     edi                    ; get command code from first byte
00160    mov     edi,edi                ; move past it
00160    mov     eax,dword ptr [Xlat_FUNCTAB[edx*4]] ; find handler
00167    xchg   eax,[esp]              ; put its address on stack
00168    retn                                ; "return" to it
V86MMGR Xlat_API    endp
```

Listing 3-14 shows V86MMGR's handler for Xlat\_API\_Exec\_Int, which reflects protected mode software interrupts to V86 mode.

#### Listing 3-14: V86MMGR Code to Handle Xlat\_API\_Exec\_Int

```

00168    push    eax
0016c    movzx   eax,byte ptr [edx]    ; get byte #2 (INT num) from script
0016f    vmmcall Begin_V86_Exec        ; set up for v86 mode
00175    vmmcall Exec_Int              ; interrupt number => in EAX
```



```
0017B pop eax
0017C VMCall End_Nest_Exec
```

The functions `Begin_Nest_V86_Exec`, `Exec_In`, and `End_Nest_Exec` are all described for the VMM section of Microsoft's DDK documentation. `Begin_Nest_V86_Exec` sets the virtual host machine to V86 mode. `Exec_In` simulates a specific interrupt in the VM, and `End_Nest_Exec` returns everything to normal. We could of course now take apart VMM and see how these functions too are implemented, but while awaiting that, it is worth taking a detour which, however, we will look at VMM a little later, to see how it implements some of the DPMI functions.

Note that Windows runs DOS in a real mode but a V86 task. VMM on a Windows VxD may trap software interrupts, IN or OUT instructions, memory access, and so on from DOS. This means that many serious DOS tasks will cause a trap back into the Windows VMM.

The simplest possible protected mode INI 21h call is enhanced mode, takes the state we've just described. DPMIGR gets the INI 21h call, figures out what it does (need to load debugging specific, and uses a script to call `VM86MMIO` to either the INI 21h call to V86 mode. `VM86MMIO` figures out what DPMIGR is asking for, and then asks VMM to switch to V86 mode, reason for INI 21h, and switch back to protected mode.

Yes, all this happens automatically. Windows program makes even the simplest INI 21h call. And we glossed over many details, such as how the PM call back actually makes the transition to protected mode, at Ring 3 to Ring 2 to protected mode at Ring 0, but we get the idea.

You might think that INI 21h calls don't interrupt any other unprotected mode. However, the DDK debug version (`VM86IN`) includes a DPMIGR debug command "Display DOS trace profile" using the name of INI 21h calls. It puts a few seconds time constraint, so it only shows a few thousand DOS calls. The most popular are INI 21h AH=20 (DOS call `VM86ExecData`), AH=10 (and `Set`), and AH=50h (Set PSP). It is considered that Windows runs these functions.

Windows likes to run in protected mode. For example, because every time Windows is an associated PSP, you won't find a task whose kernel would call INI 21h AH=50h (Set PSP) and expect switched to a given task. It has a slot at `KeBugCheckEx` (40400000) (page 546) set, and kernel puts off making the Set PSP command the task has made some other DOS call, which can be made to appear pretty frequently, even if the exception caused by a single, but extreme DPMIGR PSP (not a syn, with the Windows PSP) is a debug file, see what kernel does this, call it out, and update all the activity that occurs, and the system counter for complete INI 21h calls is `KeBugCheckEx`.

Actually, the Set PSP is not a good example of something that requires a full exit to real mode handling from DPMIGR. INI 21h AH=50h expects a PSP address in EAX. In protected mode, of course, it makes sense for this task, the protected mode selector has PSP. DPMIGR reasons the INI 21h to V86 mode, and because DOS, of course, can't handle protected mode PSPs, it must first convert the PSP in real mode segment address over to VMM. `SelectorMap` is a function.

As another example, a function for `VM86MMIO` must provide special handling, consider INI 21h AH=28h (Set Interrupt Vector). When called in protected mode, this function normally sets a protected mode interrupt vector. It is, of course, a real mode call to DOS (running in V86 mode). Instead, DPMIGR must handle the call using a transition to real mode, call to the VMM Set PM Int Vector function, and use the same function that DPMIGR used back in listing 3-12 to install its own INI 21h protected mode call back. If you actually want to install a real mode interrupt handler from protected mode, you must call down to real mode by hand, for example, by using the DPMI Simulate Real Mode Interrupt, which we will describe later.

Finally, INI 21h AH=40h (Allocate Memory Block) should in protected mode allocate extended memory, and return protected mode selectors. It would make sense in a protected mode version of this function to allocate conventional memory, and return real mode parameters, address, so again DPMIGR can call down to V86 mode. Instead, DPMIGR handles 21-40h using VMM functions such as `PageAllocate`, `AllocateEIPSelector`, and `SetSelector`. It is a bit of a

INT 21h calls its “VxD load” without calling down to V86 mode—were taken to its logical conclusion Windows wouldn't need DOS at all. This is what Chicago provides: 32 VxDs such as VFAT 386.

### How DOSMGR Handles Undocumented DOS Calls

So how does DOSMGR handle undocumented DOS functions? We already saw what it does with Set PSP calls, but this call was totally documented in DOS 5.0, so it's no longer a good example. Let's return to our old friend, Function 52h. As we'd know by heart, this function takes no parameters and returns a far pointer to SYSVAR in ES:BX. We've seen that Windows somehow transparently makes this work in protected mode, since calling INT 21h 3Ah 52h in protected mode gets back a valid protected mode pointer to SYSVAR. Well, it's not transparently sort of works, since SYSVAR itself contains real mode SYSVAR, but DOSMGR is doing the best that is humanly possible. DOSMGR can't very well modify SYSVAR to make it contain protected mode pointers, since this would immediately crash DOS.

It *can* do what DOSMGR gives us the illusion of a protected mode DOS that even supports Function 57h. The DOSMGR script that handles Function 52h looks like this:

```
Xlat_API_Return_Ptr ES, BX
Xlat_API_Exec_Int 21h
```

DOSMGR also uses this same exact macro to handle the newly documented INT 21h AH=34h (Get F-DOS File Address), which also takes no parameters and returns a far pointer in ES:BX. You can see how easy it is to fit INT 21h functions into categories and using the V86MMGR MIAI API facility make it a bit easier to implement a DOS extender.

As we'll see in my explanation for the DDK, its job is to take a real mode pointer in the specified registers from ES:BX and turn it into a protected mode pointer. The DDK notes that, although this macro is used by our MIAI API Exec Int macro in our MIAI script, V86MMGR actually creates the segment address after setting up the interrupt, showing how V86MMGR implements MIAI Return Protected mode to take care of the subject, even for this book, so we'll simply note that this is the custom of the service. As an MIAI Exec Int 21h, and then fiddles with the VME selector structure, the MIAI Return Protected turns ES and BX into Chant ES and Chant BX. For 16-bit software, the VMM Map 1 to VM Addr function (documented in the DDK) does the standard to protected mode pointer conversion.

This is certainly important and a very good example of why it is sometimes helpful to know such low-level implementation details. The DDK file, instruction for Xlat\_API\_Return\_Ptr says “For 16-bit protected mode applications that use a TSS selector, an appropriate selector does not already exist.” What it doesn't discuss is how these selectors exist got fixed. Once you know that Xlat\_API\_Return\_Ptr is implemented using Map 1 to VM Addr, you can turn to the DDK documentation for that function and see that this is precisely the “A virtual device must *never* free a selector that is not owned by the device. For this reason, this service should be used sparingly.”

It's a good thing that it actually makes a lot of sense. In the case of providing a protected mode to the user, INT 21h 3Ah 52h, how would DOSMGR know when we were done using SYSVAR and the other protected mode pointers that it created with Xlat\_API\_Return\_Ptr and thus with Map 1 to VM Addr, not be freed. There is a way to know, which is an inherent problem with transparency. If there remains a danger that this service could consume selectors, of which only 8192 are available, it would never be freed until the DOS box was far gone away, perhaps by a user who had started working around by using Map 1 to VM Addr, never to avoid this danger by keeping a lot of selectors that create a lot of trouble to help to request by returning some previously allocated selector, rather than allocating a new one.

DOSMGR uses the same Xlat\_API\_Return\_Ptr macro to handle the recently documented INT 21h functions Hb Get Default DFB and 32h Get DFB, and also to handle functions 1Bh Get Default Drive Data and 1Ch Get Drive Data. Because V86MMGR doesn't know the size of the returned





DPMI also provides a higher level function, Simulate Real Mode Interrupt, INT 31h AX=4300h, which takes care of the low level mode switching for you. We use this in effect extensively in the remainder of this chapter.

For those occasions when you need to call a real mode function with a real pointer address, DPMI also provides two Call Real Mode Procedure functions, INT 31h AX=0300h and 0302h. All of these functions are documented in the DPMI specification, available free of charge from Intel and Ray Duncan's *Extending DOS* and A. Williams' *DOS and Windows: Protection of Mod*.

While very few programs use Windows 3.0 any more, and even fewer use Windows 3.0 Standard mode, we probably ought to note that INT 31h functions 0300h, 0301h, and 0302h exist in Windows 3.0 Standard mode. According to a Microsoft article, "Using DPMI Functions Not Recreated in Std Mode" (MSK1-Q70891), this was done so that Windows 3.0 displays:

INT 31h AX=0300h is slightly confusing because you use this software interrupt to generate actually simulate you I see what it means below, in case of a 386 mode the software interrupt such as INT 21h AH=52h or AX=5100h that you are actually interested in. You use this software interrupt to simulate another.

But simulating real mode interrupts is not half the people. It is in fact not, you do not say INT 31h AX=0300h to call INT 21h AX=52h or AX=5100h in real or A86 mode, you use *real mode pointers*. The fact that you are calling the real mode software interrupt using DPMI does not change this. This protected mode programs also receive the ability to get these real mode pointers into the protected mode address space, that is, to convert real mode pointers to something called a protected mode. When we refer to these three protected mode pointers in a structured conventional memory. Sometimes we will use these structures and/or some other additional real mode pointers which also must be mapped into a protected mode program's address space.

How are real mode pointers translated to protected mode. We say "translate" but we just do, seems DOSMGR uses VxD API Return Ptr which translates Map Ptr to VM Address, can permanent selectors. This corresponds exactly to the DPMI Segment Descriptors from INT 31h AX=2, which is a dead very common one, but which Windows 3 checked code interprets as a Map Ptr To VM Addr, which is not in the selectors that can be used. For this, we need a different technique so that we can use these things when we do the work. This descriptor's only eight bytes, so this does not sound like a big deal, but the VM descriptor's 8, 9, 10 descriptors, and therefore selectors, because descriptors tables are limited to 64K. To avoid this, for reasons we won't get into here, there are only 4,096 available descriptors. I don't want to be freeing mapped real mode pointers when you're done with them, so essential Windows 3.0 will free these for you when your program terminates.

As with one selector, you'll always need protected mode pointers in the CPU space. Windows preallocates several selectors in popular memory locations. These selectors are available for use by Windows programs and drivers, these predefined selectors, whose names are 0000h,

00401h and 40001h are described and higher see a *Undocumented Windows*. Apparently as a "favor" to the Real mode Systems DOS, a number used by Gates 1, 2, 3, Windows also supports selector 40h which is a "normal" selection, the selector 40h equals to base address pointer 4000 data area at 400h. As Matt Patrick shows in Chapter 4 of his *Windows Internals*, where the Windows kernel routines creates these selectors using INT 31h AX=0002. Unfortunately, these predefined selectors only cover a portion of the real mode address space. Also, there are only a handful of Windows programs protected mode DOS programs can't get at them.

But it's still no worry, you can do about these predefined selectors because you can easily create your own protected mode selectors that map real mode addresses. This requires three DPMI functions, Allocate E81 Description, INT 31h AX=0, Set Segment Base Address, S=31h AX=7, and Set Segment Limit, INT 31h AX=8. As shown by the following code, a descriptor selector's base set its base to correspond to the real mode address. You should set the limit of the selector to be the

etc. The stack may be re-interpreted rather than blindly make everything 64k. As DOSMGR has no idea how to do it, it has no way of knowing which objects your program will want to use. Setting a small amount of stack space, using the size of the structure, or relying helps catch a lot of one errors and is better than nothing. Rather, because DPMI Windows programs can use three Windows API functions: `MemSelector`, `SetSelectorBase`, and `SetSelectorLimit`. The last two functions were at one point `FreeSelector`, see *Enhancement of Windows*. Chapter 5 examines them in detail. Having created the protected mode selector, you can turn it into a usable far pointer with a macro such as `MAK_1P` (for `land C++`) `DOS 31` or `MAK_1P` (Windows 11).

Something that we want to be able to free one of these mapped protected mode selectors when we call `FreeSel` (INI 31h AX) Free I/D Descriptor, provides this ability. The equivalent Windows API function is `FreeSelector`.

The `ANTDRIV`, `CURDRIV`, and `MEMMGR` programs printed out various addresses returned from `INt`. The addresses 52 and 54000h. The pointers that our protected-mode programs use, how we set them, is a matter of protected mode addresses created using the DPMI or Windows API functions and must be used to assign pointers to these addresses. These are far pointers, so the selector whose base address we set to the memory address we're interested in. Got it. The `sel` value itself is more or less meaning, it's a selector, a descriptor, a descriptor-like. Printing out these protected mode addresses, they will not be useful. What we'd like to do is print one of these protected mode addresses to get back its real address, the real mode address, or create the real mode address we actually want to display. This is a real mode address. DPMI Get Segment Base Address function (INI 31h AX) or in the case of Windows API function `GetSelectorBase`. Of course, only base addresses less than `0x00000000` or `0x00000000` can be used to have real mode equivalents.

Real mode, low, 16-bit, DOS functions expect us to pass them pointers, which they will fill with pointers. For example, `INt 2h` `MEMMGR` works this way to call the `realmode` function to print out real mode addresses, the real mode intercept, that that's it, it's enough. The `realmode` address, `INI 31h` `AX` `00000000`, is a real mode address. This is not to mean that the address is a real mode address, below, no, it's a pointer, or a pointer, how could DOS access it? It's not, so we can't use it to calculate the correct real mode buffer with a protected mode pointer.

Another way to get the `realmode` address, such as a transfer buffer, which `DOSMGR` accesses, is to use `INt 2h` `MEMMGR` all the data, and the `realmode` address. We can allow it, our own way, to use the `realmode` address, `INI 31h` `AX` `00000000`.

This is a real mode address, the real mode segment, paragraph, address of a real mode memory block, so we can use it to print out protected mode selectors. Programs can also write to the real mode memory block, so we can use it to print out protected mode selectors. Then pass the real mode paragraph address to `DOS` (INI 31h AX) `00000000`, `DOS` Memory Park, fixes these buffers. This is a real mode Windows API function, `GlobalDOSAlloc`, and `GlobalDOSFree`.

Something that we want to be able to document DOS from protected mode Windows, you need the following actions:

- `MemSelector` (INI 31h AX) `00000000` or `MemSelector`
- `Free I/D Descriptor` (INI 31h AX) `00000000` or `FreeSelector`
- `Get Selector Base Address` (INI 31h AX) `00000000` or `GetSelectorBase`
- `Set Segment Base Address` (INI 31h AX) `00000000` or `SetSelectorBase`
- `Set Segment Limit` (INI 31h AX) `00000000` or `SetSelectorLimit`
- `Alloc. DOS Memory Block` (INI 31h AX) `00000000` or `GlobalDOSAlloc`
- `Free DOS Memory Block` (INI 31h AX) `00000000` or `GlobalDOSFree`
- `Simulate Real Mode Interrupt` (INI 31h AX) `00000000` or `Windows API equivalent`

Now, these functions have been discussed in many other places. A number of good articles describe how to use DPMI to access real mode DOS from protected mode Windows, for those times

when the Windows DOS extender doesn't do some work for you. It may seem that we could have, as mentioned, these eight functions at the beginning of this chapter and been done with it without looking at how DOSMMAP is implemented or how Windows programmers just protected-mode DOS applications (DPMI clients) and other systems (say, so much more) working at the happenings of *data* use these DPMI functions. However, many times programmers use these DPMI functions in a cookbook fashion, without truly understanding what they actually do or why you use one of them and not another. For example, we've often seen programmers try to use INT 31h AX=0 or GetSelectorBase to call the DOS truncate function from protected mode, and waste large amounts of time because it doesn't work. In the long run, it pays to truly understand what you are doing.

### Hiding DPMI

Actually, though, having taken so long just to get to the point where most treat the 90s of 1993, we'll usually start, we're now going to spend very little time looking at these eight DPMI and only instead we're going to look them beneath a higher-level user-visible interface.

For example, to map a real-mode address space, you'd have to worry about combining all the different DPMI functions. We're going to simplify this and change it and combine with a `map_real` function that takes a real-mode pointer and the number of bytes you'd want to look at from protected mode; it returns an equivalent protected-mode pointer. When you are done using the protected-mode pointer, you can free it with `free_mapped_real`:

```
void far *prot_ptr = map_real(real_ptr, number_of_bytes),
// use prot_ptr
free_mapped_real(prot_ptr);
```

Let's continue to hide DPMI. Using our interrupt INT 31h AX=0300h to generate another such as INT 21h AH=52h is confusing. But calling a function to generate a selector (get\_sel) can bring at all DOS programmers used to do this, all the time, with functions such as `get_sel` (see Chapter 2). We're going to complement `map_real` with a function that works just the opposite: `get_sel` generates a real-mode selector, interruptible protected mode. We can use this function together with `map_real` to accomplish most of what we need. For example:

```
union REGS r;
struct SREGS s;
void far *real_sysvars,
void far *prot_sysvars,
memset(&s, 0, sizeof(s)),
r.h.eh = 0x52,
real_intBox(0x21, &r, &r, &s),
real_sysvars = PK_FP(s.es, r.x.bx),
prot_sysvars = map_real(real_sysvars, number_of_bytes),
// access SysVars via prot_sysvars
free_mapped_real(prot_sysvars);
```

These contains `real_intBox` and `map_real`. For the use of a library of routines (PROT.C) and PROT.H (see Listings 3-16 and 3-17) for protected-mode programming with Windows and another DPMI server in a Windows program, PROT.C uses Windows API functions such as `AllocSelector` and `SetSelectorBase`. This DOS-based DPMI program, PROT.H maps these to PROT.C functions such as `dpmi_alloc_sel` and `dpmi_set_selector_base`, which provide C wrappers for the appropriate DPMI INT 31h calls.

### Listing 3-16: PROT.C

```
/*
PROT.C
Routines for protected mode programming with Windows and DPMI
Andrew Schulman, February 1993
from 'Undocumented DOS', 2nd edition (Addison-Wesley, 1993)
```

```

(changed version of PROTMOSE, from RSJ, Oct 1992, Dec. 1992)
Requires 286+ instruction set (e.g., bcc -2 or cl -62)
*/
#include <stdlib.h>
#include <string.h>
#include <dos.h>
#ifdef DPMI_APP
#include "windows.h"
#endif
#include "prot.h"
/*****
static WORD mapped = 0, // to ensure don't have selector leak
WORD get_mapped(void) { return mapped; }
/*****
static DWORD base = 0; // for data in other VMs
void set_base(DWORD b) { base = b; }
DWORD get_base(void) { return base; }
/*****
void far *map_linear(DWORD lin_addr, DWORD num_bytes)
{
    WORD sel,

    /* allocate a selector similar to our current DS
    (i.e., a data selector) */
    _asm mov sel, ds
    if ((sel = AllocSelector(sel)) == 0)
        return (void far *) 0,

    /* set the base and limit of the new selector; variable "base"
    allows access to data in other VMs */
    SetSelectorBase(sel, base + lin_addr);
    SetSelectorLimit(sel, num_bytes - 1);

    mapped++;

    /* turn into a far pointer */
    return MK_FP(sel, 0);
}

void free_mapped_linear(void far *fp)
{
    freeSelector(FP_SEG(fp)),
    mapped--;
}

void far *map_real(void far *fp, DWORD size)
{
    return map_linear(MK_LIN(fp), size);
}
/*****
/* Performs a real mode interrupt from protected mode */
BOOL dpmi_rmode_intr(WORD intrno, WORD flags,
WORD copywords, RMODE_CALL far *rmode_call)
{
    // ignore flags
    _asm push di
    _asm push bx
    _asm push cx
    _asm mov ax, 0300h // simulate real mode interrupt
    _asm mov bx, intrno // interrupt number, flags
    _asm mov cx, copywords // words to copy from pmode to rmode stack
    _asm les di, rmode_call // ES:DI = address of rmode call struct

```



```

    _asm int 31h          // call DPMI
    _asm jc error
    _asm mov ax, 1      // return TRUE
    _asm jmp short done
error:
    _asm mov ax, 0     // return FALSE
done:
    _asm pop cx
    _asm pop bx
    _asm pop di
)
/*****
int real_int86(int intno, union REGS *inregs, union REGS *outregs,
struct SREGS *sregs)
(
    RMODE_CALL r,
    memset(&r, 0, sizeof(r)); // initialize all fields to zero important!
    r.edi = inregs->x.di;   r.esi = inregs->x.si,
    r.ebx = inregs->x.bx;   r.edx = inregs->x.dx,
    r.ecx = inregs->x.cx;   r.eax = inregs->x.ax,
    r.flags = inregs->x.cflag,
    r.es = sregs->es,      r.ds = sregs->ds;
    r.cs = sregs->cs,
    // NOTE: r.sp=0 so that DPMI host provides real mode stack!
    if (!dpmi_rmode_intr(intno, 0, 0, &r))
    {
        outregs->x.cflag = 1; // error: set carry flag!
        return 0;
    }
    sregs->es = r.es,      sregs->cs = r.cs;
    sregs->ss = r.ss,      sregs->ds = r.ds,
    outregs->x.ax = r.eax,  outregs->x.bx = r.ebx,
    outregs->x.cx = r.ecx,  outregs->x.dx = r.edx,
    outregs->x.si = r.esi;  outregs->x.di = r.edi,
    outregs->x.cflag = r.flags & 1; // carry flag
    return outregs->x.ax,
)
int real_int86(int intno, union REGS *inregs, union REGS *outregs)
(
    struct SREGS sregs,
    memset(&sregs, 0, sizeof(sregs)),
    return real_int86(intno, inregs, outregs, &sregs),
)
int real_intdos(union REGS *inregs, union REGS *outregs)
(
    return real_int86(0x21, inregs, outregs);
)
int real_intdosx(union REGS *inregs, union REGS *outregs,
struct SREGS *sregs)
(
    return real_int86(0x21, inregs, outregs, sregs),
)
/*****
WORD far(word wSeg) // load access rights
(
    _asm far ax, wSeg
    _asm jnz error
    _asm shr ax, 8
    _asm jmp short done; // value in AX
error:
    return 0; // can't be a valid AR
done:
)

```

```

void far *lin_to_real(DWORD lin_addr)
{
    if ((lin_addr > 0x10FFEF) /* allow MWA pointers up to FFFF:FFFF */
        return NULL; /* not accessible in real mode */
    else
    {
        WORD seg, ofs;
        seg = ((lin_addr > 0x100000L) ? 0xffff : (lin_addr >> 4));
        ofs = lin_addr - ((DWORD) seg << 4L);
        return MK_FP(seg, ofs);
    }
}

void far *get_real_addr(void far *fp) /* prot to real */
{
    DWORD base;
    if ((far)FP_SEG(fp) == 0)
        return NULL; /* not a valid pointer */
    base = GetSelectorBase(FP_SEG(fp));
    return lin_to_real(base + FP_OFF(fp));
}

//*****
BOOL alloc_real_seg(DWORD bytes, WORD *ppara, WORD *psel)
{
    DWORD dw = GlobalDosAlloc(bytes);
    if (!dw) return 0;
    *ppara = H(WORD)dw;
    *psel = (DWORD)dw;
    return 1;
}

BOOL free_real_seg(WORD sel)
{
    return (GlobalDosFree(sel) == 0);
}

//*****
#ifdef DPMI_APP
WORD dpmi_alloc_selector(void)
{
    asm xor ax, ax
    _asm mov cx, 1
    _asm int 31h
    _asm jc error
    _asm jmp short ok
error:
    return 0;
}
// return value in AX

WORD dpmi_free_selector(WORD sel)
{
    asm mov ax, 1
    _asm mov bx, sel
    _asm int 31h
    _asm jc error
    _asm jmp short ok
error:
    return sel; // act like the Window function
ok:
    return 0;
}

WORD dpmi_set_selector_base(WORD sel, DWORD base)
{
    _asm mov ax, 7

```

```

    _asm mov bx, sel
    _asm mov cx, word ptr base+2
    _asm mov dx, word ptr base
    _asm int 31h
    _asm jc error
    _asm jmp short ok
error:
    return sel,
ok:
    return 0;
)
DWORD dpmi_set_selector_limit(WORD sel, DWORD limit)
{
    _asm mov ax, 8
    _asm mov bx, sel
    _asm mov cx, word ptr limit+2
    _asm mov dx, word ptr limit
    _asm int 31h
    _asm jc error
    _asm jmp short ok
error:
    return sel,
ok:
    return 0;
)
DWORD dpmi_get_selector_base(WORD sel)
{
    _asm mov ax, 6
    _asm mov bx, sel
    _asm int 31h
    _asm jc error
    // return CX DX in DX,AX
    _asm mov ax, dx
    _asm mov dx, cx
    _asm jmp short done
error:
    return 0, // problem: indistinguishable from 1-byte segment!
done:;
}
/* C interface to LSL (Load Segment Limit) instruction */
#pragma warn -rvl
WORD lsl(WORD sel)
{
    if (!sel) return 0, /* workaround 386 bug: Hummel, p.471 */
    _asm lsl ax, sel
}
DWORD dpmi_get_selector_limit(WORD sel)
{
    // no DPMI function, use LSL, should use 32-bit LSL though!
    return (DWORD) lsl(sel);
}
DWORD dpmi_dos_alloc(DWORD bytes)
{
    DWORD retval,
    WORD paras = (bytes >> 4) + 3;
    _asm mov ax, 0100h
    _asm mov bx, paras
    _asm int 31h
    _asm jc error
    _asm mov word ptr retval, ax
    _asm mov word ptr retval+2, dx
    return retval;
error:

```

```

    return 0L,
}
WORD dpmi_dos_free(WORD sel)
{
    _asm mov ax, 0101h
    _asm mov dx, sel
    _asm int 31h
    _asm jc error
    return 0,
error:
    return sel;
}
#endif

```

### Listing 3-17: PROT.H

```

/*
PROT.H -- see PROT.C
Andrew Schulman, February 1993
from "Undocumented DOS", 2nd edition (Addison-Wesley, 1993)
*/

#ifndef _PROT_H
#define _PROT_H

#ifdef BPMI_APP
typedef int BOOL;
typedef unsigned char BYTE;
typedef unsigned short WORD;
typedef unsigned long DWORD;
#define LOWORD(dw) ((WORD)(dw))
#define HIWORD(dw) ((WORD)((DWORD)(dw) >> 16))
#else
#include "windows.h"
#endif

void far *map_real(void far *fp, DWORD size);
void set_base(DWORD b);
DWORD get_base(void);
void far *map_linear(DWORD lin_addr, DWORD num_bytes);
void free_mapped_linear(void far *fp);
WORD get_mapped(void);
void far *get_real_addr(void far *fp);
BOOL alloc_real_seg(DWORD bytes, WORD *ppara, WORD *psell);
BOOL free_real_seg(WORD sel);

#define free_mapped_real(x) free_mapped_linear(x)

#ifndef MK_FP
#define MK_FP(seg, ofs) \
    ((void far *) ((unsigned long) (seg) << 16) | (ofs))
#endif

#define MK_IN(fp) \
    (((DWORD) FP_SEG(fp) << 4) + FP_OFF(fp))

typedef struct {
    unsigned long edi, esi, ebp, reserved, ebx, edx, ecx, eax;
    unsigned flags, es, ds, fs, gs, ip, cs, sp, ss;
} KMODE_CALL;

BOOL dpmi_rmde_intr(WORD intno, WORD flags,
    WORD copywords, KMODE_CALL far *rmde_call);

int real_vec86x(int intno, union REGS *inregs, union REGS *outregs,
    struct SREGS *sregs);

int real_int86(int intno, union REGS *inregs, union REGS *outregs);
int real_intdos(union REGS *inregs, union REGS *outregs);
int real_intdosx(union REGS *inregs, union REGS *outregs);

```

```

    struct SREGS *sregs);
// to be defined by application;
// if DPMM_APP, app's fail() must call _dos_exit'
extern void fail(const char *s, ...);
#ifdef DPMM_APP
#define AllocSelector(x)      dpmm_alloc_selector()
#define FreeSelector(x)      dpmm_free_selector(x)
#define SetSelectorBase(x,y) dpmm_set_selector_base((x), (y))
#define SetSelectorLimit(x,y) dpmm_set_selector_limit((x), (y))
#define GetSelectorBase(x)   dpmm_get_selector_base(x)
#define GetSelectorLimit(x)  dpmm_get_selector_limit(x)
#define GlobalDosAlloc(x)    dpmm_dos_alloc(x)
#define GlobalDosFree(x)     dpmm_dos_free(x)
WORD dpmm_alloc_selector(void);
WORD dpmm_free_selector(WORD sel);
WORD dpmm_set_selector_base(WORD sel, DWORD base);
WORD dpmm_set_selector_limit(WORD sel, DWORD limit);
DWORD dpmm_get_selector_base(WORD sel);
DWORD dpmm_get_selector_limit(WORD sel);
DWORD dpmm_dos_alloc(DWORD bytes);
WORD dpmm_dos_free(WORD sel);
#endif
#endif

```

In addition to real `int86x`, there are the real `int86`, real `intdos`, and real `intdosx` functions. All of these functions in turn call `dpmm_mode_int` and do little more than convert from the common `RREGS` and `struct SREGS` structures. `DPMM` (see Chapter 2) is `DPMM` specifies `RMODE_CALL` structure. The `dpmm_mode_int` function also `PROTECT` provides a wrapper for the `DPMM` `Real Mode` interrupt function `INT 31h AX 0500h`. Of course, you can call `dpmm_mode_int` directly if you prefer this to real `int86x`.

About the only thing worthy of note is that real `int86x` zeros out all fields in the `DPMM` `RMODE_CALL` structure and in particular forces the `SSP` fields in the structure to zero. It tells the `DPMM` host to provide a real mode stack rather than expecting us to provide one. So we use that 16 segment register fields in the structure should contain values that are near zero. But real `int86x` is not `protect` for a protected mode program to initialize these fields with the values of the actual segment registers from `segment` for example because these are protected mode values which real mode would interpret incorrectly.

The `map_real` function is more interesting. As noted earlier, this function hides the work of three `DPMM` functions. Actually `map_linear` is the function that does all the work; `map_real` consists of just a single call to `map_linear` using the `MK_FP` macros from `PROTECT`. Improving the familiar rule of real mode addressing `linear = seg * 10h + offset`, `MK_FP` converts a segment/offset real mode address to so-called *linear address*. A linear address is an indication of the protected mode address space. If paging is enabled, as is much the case in Windows, this address mode, then it is not necessarily a physical address. This is an important issue, but one that need not concern us right now.

The `map_linear` function calls `INT 31h AX 4` or `AllocSelector` to allocate a new protected mode selector; the function also calls `INT 31h AX 7` or `SetSelectorBase` to set the new selector's base address to the linear address parameter. Optionally, `map_linear` can add a base of a word we'll use this later (see Listing 3-33 for example) for accessing data on other virtual machines. By default the base is zero, so `map_linear` creates selectors that map data to the current VM. The function next calls `INT 31h AX 8` or `SetSelectorLimit` to set the new selector's limit to the size that was requested, actually to call less than the size requested since the limit is not the size of the segment but its last valid byte offset (a one byte segment has a limit of zero, for example). Finally, it, to, that uses the `MK_FP` macros to turn this selector into a far pointer.

The resulting far pointer is just as usable in protected mode as regular far pointers are in real mode. For example, to access four bytes at (and at address 12345h) in real mode you could just create a pointer like `0x0000:0000+12340005` or `12340005` or any number of other segment/offset combinations. To access this 4 bytes in protected mode, you need to create a map (near `0x12345:4`) this returns a segment selector, whose offset is zero and whose segment is some essentially meaningless value, as indicated in descriptor table. But the segment is underling base address will be 12345h. However, the far pointer, when we create far pointers all we really care about is getting to some memory address, so we can calculate it using the pointer's segment and offset. If the pointer ends at offset 5, we can write to `12345h` it could be `0000:0000` for all we care. Think of map real as a stand-in for `0x0000:0000` for real mode pointers.

The only thing that can do is incrementally count of mapped selectors (free mapped linear addresses) that you get on requests the counter with `get_mapped_linear_selector` really does all the necessary system calls, but is not automatically deallocated when your program exits. It is true that we have a possibility selection we allocate. By checking that `get_mapped_linear` is zero just before your program terminates you can ensure that you have been a good citizen.

### Fixing SFTWALK

We can use the same DOS stand-alone routines to redo our earlier applications. It would be tedious to go back and look for errors if they were, I just redo SFTWALK.C from Listing 3-5 and a new WIN\_APP.C and C:\PROJ\WIN\_APP.C as sources for the reader. Listing 3-18 shows the new version of SFTWALK.C. Its program can be linked with a static DPMI\_APP.LIB containing the DPMI.C, R.C, and EBC.C codes (as shown earlier, Listings 3-13, 3-14, and 3-16). In addition to the DPMI\_APP.LIB, SFTWALK now also needs a library for making a Windows version. The Windows version is available as a library symbol by linking with WIN\_APP.LIB, which contains the PRINTF.LIB, `WIN_APP.LIB` (PROJ) module. WIN\_APP.LIB, of course, does not include the DPMI module.

### Listing 3-18: SFTWALK.C (Fixed for Protected Mode)

```

/*
SFTWALK.C -- Count FILES by Walking SFTs
Version #2, uses real, int80() and map_real() in PROT.C
Andrew Schulman, March 1993
From "Annotated DOS", 2nd edition (Addison-Wesley, 1993)

DPMI program
b.c -c DPMI_APP ? dpmish.c printf.c ctrl.c asm
f -to dpm_app -dpmish obj -ctrl.c obj -prot.obj
br -c DPMI_APP ? sftwalk.c dpmi_app.lib

Windows program
bcc -c -W -2 -BWINDOWS printf.c prot.c
Lib win_app ->printf.obj ->prot.obj
bcc -W -2 -BWINDOWS sftwalk.c win_app.lib
*/

#include <stdlib.h>
#include <string.h>
#include <dos.h>

#define DPMI_APP
#define PROT_MODE
#include "dpmish.h"
#include "prot.h"
#endif

#define WINDOWS
#define PROT_MODE
#include "windows.h"
#include "prot.h"
#include "printf.h"

```

```

#else
#include <stdio.h>
#endif

typedef unsigned char BYTE,
typedef unsigned short WORD,

// moved here (and changed) from SYSVARS.C
BYTE far *get_sysvars(void)
{
#ifdef PROT_MODE
union REGS r;
struct SREGS s;
r.h.eh = 0x52,
/* do not pass in garbage for seg regs, but don't use
segread() either! Use memset to initialize to zero */
memset(&s, 0, sizeof(s));
s.es = r.x.bx = 0;
real_int26x(0x2f, &r, &r, &s);
return (BYTE far *) map_real(RK_FP(s.es, r.x.bx), 0xffff),
#else
_asm xor bx, bx
_asm mov es, bx
_asm mov eh, 52h
_asm int 2fh
_asm mov dx, es
_asm mov ax, bx
// return value in DX:AX
#endif
}

typedef struct _sft {
struct _sft far *next;
WORD num,
// other stuff not used here
} SFT;

int do_sftwalk(void)
{
BYTE far *sysvars,
SFT far *sft, far *next,
WORD num_mapped,
int files = 0;

if (! (sysvars = get_sysvars()))
return 0;
sft = *((SFT far *) &sysvars[4]);
next = sft;

#ifdef PROT_MODE
sft = (SFT far *) map_real(sft, sizeof(SFT));
#endif

while (FP_OFF(sft) != 0xffff)
{
files += sft->num;
#ifdef PROT_MODE
{
// print out, using saved real mode address
// (printing out the protected mode address isn't helpful)
printf("SFT @ %p — %u files\n", next, sft->num),
next = sft->next; // save ptr to next
free_mapped_real(sft); // zap old one
if (FP_OFF(next) == 0xffff) // at end of list
break;
sft = (SFT far *) map_real(next, sizeof(SFT)); // map new one
}
#endif
}
return files;
}
#endif

```

```

    printf("SFT @ %p -- %u files\n", sft, sft->num),
    sft = sft->next;
#endif
}
printf("FILES=%d\n", files),

#ifdef PROT_MODE
free_mapped_real(sysvars);
if ((num_mapped = get_mapped()) != 0)
    printf("Didn't free all selectors! %u remaining\n", num_mapped);

return num_mapped,
#else
return 0,
#endif

#ifdef DPW1_APP
void foo (const char *, ...) { puts(s), dos_exit(1), }
real_main(int argc, char *argv[]) { return 0; }
#ifdef _cplusplus
extern "C" int arg, char *argv[] { return do_sftwalk(), }
#else
// Windows program
#ifdef _cplusplus
extern "C" int PASCAL WinMain(HANDLE hInstance, HANDLE hPrevInstance,
LPSTR lpszCmdLine, int nCmdShow),
#endif
#endif
int retval;
open_display("SFTWALK ");
retval = do_sftwalk(1,
show_display());
return retval,
}
#endif

```

Note that `SEFWALK` not only calls `map_real` for `sysvars`, but also calls the function for each `SEI` pointer. If `z` of `SEI` pointer in `sysvars` is a real mode pointer, as is each next `SEI` pointer in the `real_list`. Whereas the real mode version of `SEFWALK` can follow the list with a simple `SEI = SEI->next` pointer and mode conversion gets considerably more trouble. Note that while it would be nice to have a `SEI` pointer to `SEI` in `sysvars`, this code would be incorrect because it would convert `SEI` to real mode, register it with a selector. We must stick to a call to `map_real` to determine, before passing the `SEI` pointer to `map_real`, however we must first convert `SEI` to real mode. Walking real mode linked lists from protected mode is a little tricky, but we get around it by recognizing that `SEFWALK` calls `get_mapped` before coming and will convert `SEI` to real mode any outstanding allocated selectors.

When the program `SEFWALK` now works in protected mode. It doesn't complain, so we are freeing all allocated selectors. And it produces the same output as the real mode version. Actually, that `SEI` is not the same as the `SEI` pointer that `SEFWALK` uses would not be helpful, because it is a protected mode value and gives us no indication of where the `SEI` actually is. In fact, displaying the protected mode `SEI` pointers merely implies the same address over and over because as soon as they go back to real mode, it indicates a new one. Usually the program ends up reusing the same selector number. `SEFWALK` prints out the real mode `SEI` pointer value *before* converting the pointer to protected mode with `map_real`. Remember the point made earlier that there's nothing wrong with displaying real mode pointer values in protected mode; you just can't dereference them.



SETTABLE could instead have used the `get_real_addr` function from `PRO_C` to retrieve the underlying real mode pointer from one of the protected mode pointers. However, the resulting real mode pointer is not necessarily in its most familiar form. Many real mode segment offset combinations converge to the same linear address, so the real-to-linear function is not reversible. One linear address can result in many different segment offset combinations, and the one that `get_real_addr` generates may not be the one that was fed into `map_real`. They do have the same linear address, of course.

### Inside the DPMI Server in VMM

The underlying DPMI or Windows functions used in `real_386bx` and `map_real` are such that well hidden that we can now forget about them and proceed with writing DPMI or Windows programs that use undocumented DOS.

But since we are so dependent on these underlying DPMI functions, it pays to ask what in fact they actually do, and what part of Windows provides them. Who is hooking `INT_31h` in protected mode and `INT_2Fh` in V86 mode, and what do they do when one of these DPMI functions gets called?

In standard mode, the DPMI server is part of `DOSSAVE` along with the `DOS EXEC` and `DOSV` (a bootstrapping program in its own right) general purpose DOS extender and DPMI server not really tied to Windows. But our focus here is on Enhanced mode, since that is what the vast majority of Windows users are running. The Enhanced mode DPMI server is part of VMM which you may recall is the Virtual Machine Manager upon which all VxDs including `IO$SMGR` are the native based VMM's the first 32-bit component in `WIN386-EM`. An examination of a VMM could be the subject of a fine book, especially since, with Microsoft's forthcoming Chicago, this topic will grow in importance. But for now we'll look only at the DPMI server inside VMM. What really goes on when we make a DPMI call in Enhanced mode.

First VMM hooks `INT_2Fh` in both V86 and protected mode. The `INT_2Fh` hook is used not only to handle DPMI calls such as `Function 1686h` (Get CPU Mode, used in `dpmi_present` in Listing 3-6) and `Function 1687h` (Open Real-to-Protected Mode Switch Entry Point, used in `dpmi_nfs` in Listing 3-6) but also to handle non-DPMI `INT_2Fh` calls. The calls consist of the only AP that Windows provides to non-Windows applications. Microsoft calls its most of these calls but unfortunately in a very convoluted way (see Appendix C "Windows Interrupt 2Fh Services and Notifications" in the Windows 3.1 DDK *Device Driver Adaptation Guide*).

A number of these functions are just two layers on top of VMM. For example,

- `INT_2Fh AX:1680h` (Resume Current VM Time Slice, documented in the *MS-DOS Programmer's Reference* as the `MSDOS` file Call does little more than call the VMM `Resume Time Slice` function.
- `INT_2Fh AX:1681h` (Begin Critical Section, calls VMM `Begin Critical Section`).
- `INT_2Fh AX:1682h` (End Critical Section, calls VMM `End Critical Section`).
- `INT_2Fh AX:1683h` (Get Current Virtual Machine ID, extracts the VM ID from the VMCB control block, calls the current VMM. For the VMCB, see `VMMWALK.H` in Listing 3-24 later in this chapter).
- `INT_2Fh AX:1686h` (Get CPU Mode, checks the `VMMIAE.PM_EXFC` bit in the current VMCB's status field).

VMM makes no distinction between the DPMI and non-DPMI `INT_2Fh` calls. From looking at the VMM code, DPMI appears to have started simply as a way for VMM to export some of its functionality to user privileged clients; only later did it become a separate specification shaped by vendors outside Microsoft.

We are more interested in the DPMI INT 31h functions. To install a protected mode INt 31h handler, VMM calls Allocate PM Call Back and Set PM Int Vector (the DDK documents all these VMM functions, as does Daniel Norton's book *Writing Windows Device Drivers*). The main INT 31h handler splits a few lines of code using a jump table to get to handlers for the specific DPMI function groups: AH=0 for I/O management services, AH=1 for DOS memory management services, AH=2 for interrupt management services, AH=3 for translation services, and so on. These handlers in turn invoke other handlers for the specific DPMI functions: AX=0 for Allocate LDT Descriptor, AX=1 for Free LDT Descriptor, and so on.

For example, the DPMI Segment to Descriptor function (INT 31h AX=2) was described earlier as nothing more than a layer on top of the VMM Map Lin To VM Addr function. Here's what the VMM code for this function actually looks like:

```

DAB97  DPMI_0002                ; Segment to Descriptor
DAB97  movzx eax, [ebp,Client_Rx] ; BX=real mode segment address
DAB9B  shl  eax,4                ; EAX=linear address now
DAB9E  mov  ecx,0FFFFh          ; ECX=limit (make 64k)
DABA3  vmmcall Map_Lin_to_VM_Addr
DABA9  jc  DPMI_ERROR
DABAF  mov  [ebp,Client_Ax],cx   ; return selector to caller in AX
DABB3  jmp DPMI_OK

```

Unfortunately, the DPMI function we're most interested in, Simulate Real Mode Interrupt (INT 31h AX=0300h) is sufficiently complicated that to show its implementation in full detail would require more space than we have here. We can summarize the function's operation, however, by noting that it makes the following sequence of VMM calls:

```

Push_Client_State
Begin_Nest_V86_Exec
Simulate_Int (or Exec_Int)
Resume_Exec
End_Nest_Exec
Pop_Client_State

```

This looks a lot like what Mm API Exec does when DOSMMGR asks it to reflect INT 21h to V86 mode (see Listing 3.15). In fact these functions are a lot alike. Remember, by using DPMI you are not doing in-hand what DOSMMGR or another DOS extender does or would do transparently. In essence, the DPMI Simulate Real Mode Interrupt function sets up the VM for V86 mode and tells VMM to simulate a system interrupt. Simulate means that, rather than use an actual INt instruction, VMM gets the appropriate interrupt handler's address from the low-level interrupt vector table and shuffles the stack so that when the VM runs in Resume\_Exec, it returns to this address.

The three DPMI functions used to build map linear (and map real), AllocSelector (INt 31h AX=0), SetSelectorBase (INt 31h AX=8) and SetSelectorLimit (INt 31h AX=R) are easier to describe. Although again it would be tedious to show the actual VMM code, the AllocSelector call is based on the VMM Allocate LDT Selector function (documented in the DDK). This call also uses the GetDescriptor (DWORDS) function. The two SetSelectorX calls are built using GetDescriptor and SetDescriptor.

Interestingly, the equivalent Windows API functions are *not* implemented using the DPMI calls. Instead, these functions hang directly on the LDT. *Undocumented Windows* (Chapter 5) shows the code for these two functions. According to Matt Pietrek's *Windows Internals* (Chapter 2), the Windows kernel bypasses DPMI for performance reasons. A version of the Windows KERNEL that purely used DPMI was apparently used at one time, but Microsoft found it to be too slow (and, perhaps, too portable to other operating systems). Of course, the rest of us are still supposed to use DPMI instead of hanging directly on descriptor tables.

In any case, the code *undocumented Windows* shows that these functions *can* be simple. You could, if you cared to implement your own `SelectIO` and `GetSectorIO` functions, without using Windows or DPML, by directly manipulating the FDI. There is no reason to not do this, but here, at least, there isn't much mystery about what is happening behind the scenes.

Basically, by using the DPML functions used to implement `map real`, `area` and `convenience`, the DPML *Simple Real-Mode Interrupt* call, however, carries out a lot of seemingly magical tasks that we probably could never carry out on our own.

### Back to Windows Programming

Having used `map real` and `real to80x` in DPML programs, we can now return to Windows programming and use these same two functions in a Windows application, *map real*, *use the Windows API functions* rather than the DPML IN1-314 functions, *real to80x*, *convenience*, *use the DPML function*, which has no Windows API equivalent. Windows applications *can* use DPML functions because, as noted before, Windows applications are just DPML clients. They are essentially the same as the DOS programs we dealt with in *DMISH*, except that someone—the Windows kernel—has already called IN1-210-AX-1687h and switched into protected mode for them. And they have a rather user interface, *use use a different EMU file format*, and under *undocumented real-mode programming* in a single virtual machine, called the *System VM*, rather than to separate DOS boxes—and cost a lot more to develop,” adds one tech reviewer.

Okay, so there are a lot of differences between Windows programs and protected mode DOS programs. But *everything is possible* in a virtual system. Windows applications are nothing more than fancy-looking DPML clients, and everything we learned about using DPML in protected mode DOS programs applies to Windows applications. The `cs to80x` and `map real` functions can help turn most undocumented DOS utilities into bona fide Windows applications.

### Windows and the SFT

Back in Listing 3-18, *SFTWALK* was built as a DPML program using DPML APP FIB WIN APP FIB *can now be* *run* *SFTWALK* into a Windows program, *use that* *do an I/O* *and that* *access* *determines the DOS FILES* *call* *just as the real mode and DPML versions did*. Figure 3-5 shows the Windows version of *SFTWALK*.

**Figure 3-5: SFTWALK as a Windows program**

```
SFT @ 0116:00CC -- 5 files
SFT @ 05EB:0000 -- 40 files
SFT @ 23C6:0000 -- 82 files
FILES=127
```

This is certainly a nice contrast to the GP flag message in Figure 3-3, but note that this output is different from that produced by the DPML APP session or, in the real mode version, when run outside Windows under *plain vanilla* DOS. We get three different results from *SFTWALK*:

	MS-DOS	Windows DOS box	Windows application
SFT #1	0116 00CC (5)	0116 00CC (5)	0116 00CC (5)
SFT #2	05EB 0000 (40)	05EB 0000 (40)	05EB 0000 (40)
SFT #3	none	1295 0004 (10)	23C6 0000 (82)
FILES#	45	55	127

What has happened? The differences shown in the table depend not on the type of application, but rather on *where* the application is run. For DOS boxes, VMs other than the System VM, Windows Enhanced mode has a `PerVMInfo` setting whose default is 10. *DCSMGR* adds an extra result SFT1 to each DOS box. Note that *DCSMGR* can do this only if the DOS SHARE registry is not loaded. *SHARE* must be able to communicate all open files, and it wouldn't be able to if DOS boxes had their

own private SFLs. In the System VM, the Windows kernel tries to add a much larger extra SFL to use as a free "handle cache." Chapter 4 discussed this topic in the sections "CON CON CON CON CON CON" and "But That *crashes*," and "KRN1 386 Grows the SFL." SEIWAJK lets us see the results of Windows' manipulation of the SFL chain.

When a Windows VM adds its own private SFL, Windows must instance the "next" pointer at the end of the pre-Windows SFL chain. In the examples we've been looking at here, this pointer would be the DWORD at offset 0000. Unlike a DOS internal data structure such as the CDS, the entire SFL is not instance'd; only this link from the global SFL to the private ones is. See "Instance'd Pre-Windows SFLs in Enhanced Mode Windows," "MSKB Q90796" and "Limits on the Number of Open Files," "MSKB Q81877." The SFL link instance'd, as in `DCSMGR` looks like this:

```
06377 mov eax, SFL_PTR ; end of pre-windows SFL chain
0637C mov dword ptr [esi+8],eax ; esi is inst struct
0637F mov dword ptr [esi+0Ch],4 ; 4 bytes = SFL ptr link
06386 mov dword ptr [esi+10h],200h ; instance flag = ALWAYS_Field
0638D push 0
0638F push esi
06390 VMCall _AddInstanceItem
```

As you can see, Windows goes to such trouble over the SFL. Modifying SEIWAJK to print out the actual file names in each SFL or printing the FILES program from Chapter 8, Listing 8-19, shows that the number of Windows applications you run at the same time, only a small portion of the original SFLs. In the `SYSTEMINFO` tool, section has a `CachedFiles` setting, which defaults to 2. That is, in this number appears to have no effect at least in Windows 4.1. Generally the SFL contains the names of only the files used by the most recently opened large application. This cache is supposed to help with Windows' dynamic linking; see the `FlushCachedFileHandle` description in Chapter 5 of *Undocumented Windows*.

### Walking the Device Chain

In the same way that SEIWAJK follows the linked list of SFLs from protected mode Windows, we can write a protected Windows program that walks the DOS device chain. You can build `DEV.C` (shown in Listing 3-19) in the `tools\work\DOS` or protected mode Windows. When built for Windows, the program, of course, uses the `realmode` and `mapreal` functions from `PROTC`. However, to reduce the number of built `WINDOWS` preprocessor directives, `DEV.C` employs a number of macros such as `MAP` and `FREE_MAP`, which call the appropriate functions under Windows but which do nothing under DOS. Figure 3-8 shows output from `DEV.C`.

### Listing 3-19: DEV.C for Windows or DOS

```
/*
DEV.C display MS-DOS device chain -- for Windows
Andrew Schulman, March 1993
*/
Real mode DOS: bcc dev.c

Windows
bcc -t -2 -DWINDOWS -W prot.c printf.c
t lib win_app +prot.c +printf.c
bcc -2 -DWINDOWS -W dev.c win_app.lib
*/

#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <dos.h>
#ifdef WINDOWS
#include <windows.h>
#include "prot.h"
#include "printf.h"
#endif
```

```

#endif
#define WINDOWS
char *app = "Walk DOS Device Chain";
#define puts(s)      MessageBox(NULL, s, app, MB_OK)
#endif
// this fail() not suitable for a DPMI_APP (which would
// need to do a _dos_exit()
void fail(const char *s, ...) { puts(s), exit(1), }
#ifdef MK_FP
#define MK_FP(seg, ofs) \
    ((void far *) (((unsigned long) (seg) << 16) | (ofs)))
#endif
#ifdef WINDOWS
#define MAP(ptr, bytes)      map_real(ptr), (bytes)
#define FREE_MAP(ptr)      free_mapped_linear(ptr)
#define GET_REAL(ptr)      get_real_addr(ptr)
#define YIELD()            yield()
#else
#define MAP(ptr, bytes)      (ptr)
#define FREE_MAP(ptr)      /**/
#define GET_REAL(ptr)      (ptr)
#define YIELD()            /**/
#endif
/* some device attribute bits */
#define CHAR_DEV    (1 << 15)
#define INT29      (1 << 4)
#define IS_CLOCK   (1 << 3)
#define IS_MUL     (1 << 2)
#pragma pack(1)
typedef struct DeviceDriver {
    struct DeviceDriver far *next;
    unsigned attr;
    unsigned strategy;
    unsigned intr;
    union {
        unsigned char name[8],
            unsigned char blk_cnt;
    } u;
} DeviceDriver;
typedef struct {
    unsigned char misc[8],
        DeviceDriver far *clock,
        DeviceDriver far *con,
    unsigned char misc2[18];
    DeviceDriver nu,, /* not a pointer */
    // ..
} ListOfLists; // DOS 3.1+
ListOfLists far *get_dosList(void)
{
    union REGS r;
    struct SREGS s;
    memset(&r, 0, sizeof(r));
    memset(&s, 0, sizeof(s));
    r.x.ax = 0x5200,
#ifdef WINDOWS
    /* If Windows, call undocumented DOS INT 21h function 52h via DPMI
    "Simulate Real Mode interrupt call (INT 31h AX=0300h), and return
    the resulting real mode pointer */
    real_int86x(0x21, &r, &r, &s);
#else
    int86x(0x21, &r, &r, &s),

```

```

#endif
    return (ListOfLists for *) WK_FP(s.es, r.x.bx);
}

#ifdef WINDOWS
#ifdef __cplusplus
extern "C" int PASCAL WinMain(HANDLE hInstance, HANDLE hPrevInstance,
    LPSTR pszCmdLine, int nCmdShow);
#endif
int PASCAL WinMain(HANDLE hInstance, HANDLE hPrevInstance,
    LPSTR pszCmdLine, int nCmdShow)
{
    int argc;
    char *argv[];

    ListOfLists for *doslist;
    DeviceDriver for *dd;
    DeviceDriver for *next;

#ifdef WINDOWS
    if (!(GetWinFlags() & WF_PRODE))
        fail("This program requires Windows Standard or Enhanced mode");

    // could also do protected mode DOS version with BPW1
#endif

    if (!(doslist = get_doslist()))
        fail("INT 21h Function 52h not supported");

#ifdef WINDOWS
    /* get protected mode pointer to DOS internal variable table */
    doslist = (ListOfLists for *) map_real(doslist, sizeof(ListOfLists));
    open_displayapp;
#endif

#define PARANOID
#ifdef PARANOID
    /* This block of code just double-checks that everything is ok */
    /* NULL is part of DOSLIST, not a pointer, so don't need to map */
    if (fmemcmp(doslist->nul.u.name, "NULL", 8) != 0)
        fail("NULL name wrong");
    if (!(doslist->nul.attr & IS_NULL))
        fail("NULL attr wrong");

    /* CON is pointer, so need to map */
    dd = (DeviceDriver for *) MAP(doslist->con, sizeof(DeviceDriver));
    if (fmemcmp(dd->u.name, "CON", 8) != 0)
        fail("CON name wrong");
    if (!(dd->attr & CHAR_DEV))
        fail("CON attr wrong");
    FREE_MAP(dd);

    /* CLOCKS is also pointer, so need to map */
    dd = (DeviceDriver for *) MAP(doslist->clock, sizeof(DeviceDriver));
    if (fmemcmp(dd->u.name, "CLOCKS", 8) != 0)
        fail("CLOCKS name wrong");
    if (!(dd->attr & IS_CLOCK))
        fail("CLOCKS attr wrong");
    FREE_MAP(dd);
#endif
#endif /*PARANOID*/

    /*
    Print out device chain- thanks to the MAP, FREE_MAP, GET_REAL,
    and YIELD macros, this works in real mode DOS or in protected mode
    Windows. Old real mode only code looked like this:
    do {

```

```

        printf("%Fp\t", dd),
        if (dd->attr & CHAR_DEV)
            printf("X.BFs\n", dd->u.name);
        else
            printf("Block dev: %u unit(s)\n", dd->u.blk_cnt);
        dd = dd->next,
    } while (FP_OFF(dd->next) != 1),
*/
for (dd = &doslist->nul, )
{
    printf(" %Fp\t", GET_REAL(dd)), /* print real mode addr */
    if (dd->attr & CHAR_DEV)
    {
#ifdef __BORLANDC__
        /* Borland C++ printf can't handle non-terminated strings
        char buf[9],
        _fmemcpy(buf, dd->u.name, 8);
        buf[8] = '\0',
        printf("%s\n", buf);
#else
        printf("X.BFs\n", dd->u.name),
#endif
    }
    else
        printf("Block dev. %u unit(s)\n", dd->u.blk_cnt),
        next = dd->next; /* get next pointer */
    /* first time through, this will free selector to doslist */
    FREE_MAP(dd), /* THEN free rmode seg */
    if (FP_OFF(next) == 0xFFFF) /* is there a next? */
        break,
    dd = (DeviceDriver *ar *) MAP(next, sizeof(DeviceDriver)),
    YIELD(), /* no message loop in this program, so yield */
}

#ifdef WINDOWS
{
    WORD mapped,
    if (mapped = get_mapped())
        printf("Error! %u remaining mapped selectors\n", mapped),
        show_display(),
        return mapped; /* 0 indicates success */
}
#else
    return 0;
#endif
}

```

Figure 3-6: DEV Under Windows

```

011A.0008    M:,
1A51.0002    Block dev  1 unit(s)
1A52.0004    Block dev  6 unit(s)
1A53.0006    Block dev: 3 unit(s)
CC2C.0000    CBLSSYS$
1328.0000    MSC0001
C884.0000    MSSMOLSE
0200.0000    Block dev: 1 unit(s)
0298.0002    386MAXSS
029A.0000    2MRXXXIO
026E.0000    SETVERXX
0255.0000    Block dev: 6 unit(s)
0072.0003    COM
0073.0005    AUX
0074.0007    PRN
0075.0009    CLOCKS

```

```

0076:0008      Block dev: 3 unit(s)
0077:0008      COM1
0078:0008      LPT1
0079:000F      LPT2
007B:0008      LPT3
007C:000A      COM2
007D:000C      COM3
007E:000E      COM4
7B3E:0000      @###xxxx0
    
```

An item worthy of note here is the second EMMXXXX0 expanded memory device at the very end of the list. As noted in Chapter 1 (see "Implementing DOSMMGR Functions"), VR6MMGR uses the DOSMMGR\_AskDevice function to add its emulated EMM driver onto the end of the DOS device chain.

### Truename

So far, as we've seen calling undocumented DOS functions from Windows with real-mode pointers, and mapping a protected-mode pointer with map\_real(). This works fine for functions such as INT\_21h API 52h, but it isn't always the appropriate way to call an undocumented DOS function from protected mode.

For example, the undocumented DOS Truename function (INT\_21h API-60h) expects a pointer to a ANSI path name in DS:SI; this function places the "canonical" fully qualified version of the path name in a buffer pointed at by ES:DI. Windows does not support this function in protected mode, so you must use a real-mode or some other derivative of INT\_31h API 0400h to make the INT\_21h API 60h call from real- or 386 mode. But when setting up the registers for the real-mode call, what values do you put in DS:SI and ES:DI? You obviously can't use protected-mode pointers because the whole point of using DPMI function 0400h is to generate a real- or 386-mode interrupt. A protected-mode pointer would be misinterpreted by DOS as a real-mode pointer. Obviously, then, you must put real-mode pointers in DS:SI and ES:DI. But where do you get these real-mode pointers? A moment's consideration will reveal that our old standards map\_real() and get\_real\_addr() are not useful here.

Consider how a true\_name() function would work if there were a nice C interface to DOS:

```

char *src = "c:\\this\\is\\some\\path",
char dest[72],
if (truename(src, dest))
    printf("Truename of %s is %s\n", src, dest),
    
```

in real-mode DOS, the truename() function could be implemented more or less like this (see Chapter R for a better version):

```

char far *truename(char far *src, char far *dest)
{
    __asm push ds
    __asm push d
    __asm push si
    __asm les di, dest
    __asm lds si, src
    __asm mov ah, 60h
    __asm int 21h
    __asm pop si
    __asm pop di
    __asm pop ds
    __asm jc error
    return d,
error:
    return (char far *) 0;
}
    
```



In protected mode, we want fragments to work the same way. Originally, the function will somehow have to turn the src and dest parameters into real-mode pointers. How? In a protected-mode program, src and dest will almost certainly be located in extended memory.

But there isn't any real-mode equivalent to a protected-mode pointer whose selection has a base address in extended memory. Thus it seems there is no way for fragments to readily convert its protected-mode parameters to real-mode for use by DOS. In fact, there is not. What fragments must do is take the ASM HZ string to which src points and copy it into a buffer located in conventional memory. Fragments must pass the real-mode address of this buffer to DOS. And it needs to create another buffer into which DOS can copy fragments' DOS results before returning arithmetic. First again copy this data from the conventional memory results buffer to the caller's dest pointer.

The real question then is how to allocate conventional memory from a protected-mode program. There are numerous ways to do this. For example, you could use real-mode to call the DOS allocation function INT 21h AH 48h. Calling INT 21h AH 48h in protected mode generally allocates extended memory, and certainly returns protected-mode selectors, but calling the function from real-mode without real-mode would of course void the real-mode paragraph address of a block of conventional memory. You could use some way to get a protected-mode selector to this conventional memory. You would manipulate this conventional memory block with a protected-mode address and pass the original real-mode paragraph address to DOS. The conventional memory block then acts as a transfer buffer.

Fortunately there is a easier way to allocate conventional memory from protected mode. The DPMI Allocate DOS Memory block function (INT 31h AX 0100h) allocates a conventional memory block and returns both the real-mode paragraph address and protected-mode selector pieces that you need to treat the block as a transfer buffer. DPMI also provides Free DOS Memory Block (INT 31h AX 0101h) to release it, of course. The Windows API provides two similar functions, AllocDOSBlock and GetFreeDOSBlock, of course. Here's nothing magical about these functions. This does not allow using conventional memory directly, for example. Allocate DOS Memory Block simply makes use of the DOS function INT 21h AH 48h.

PROJCTWING 3.16 provides an interesting set of functions to explore, most notably the DPMI or Windows DOS Allocation Windows Explorer to show how to use this function. Then to explore the TRUENAME.C program in Listing 3-20, contains a protected-mode filename function and is a Windows testbed for trying it out. Simply run TRUENAME with a command line argument as Windows programs can, or a command line as if TRUENAME displays a message box with the results.

### Listing 3-20: TRUENAME.C for Windows

```
/*
TRUENAME.C -- for Windows
Andrew Schulman, February 1993
bcc -W -DWINDOWS 2 truename.c win_app.lib
*/

#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <ctype.h>
#include <dos.h>
#include "windows.h"
#include "printf.h"
#include "prot.h"
#define DEST_OFS 128
#define SRC_OFS 0
void fail(char *s) { MessageBox(NULL, s, "TRUENAME", MB_OK), exit(1); }
/* see real mode version in DISKSTUF.C (Listing 8-4) */
char far *truenam(char far *src, char far *dest)
```

```

PRODE_CALL r;
char far *s2, far *fp,
WORD para, sel,

/* INT 21h AH=00h doesn't like trailing or leading blanks */
while (!isspace(*src)) src++;
s2 = src;
while (*s2) s2++;
s2--;
while (!isspace(*s2)) *s2-- = '\0';

/* Alloc 256 bytes conventional memory first 128 to transfer in src;
second half to transfer back dest */
if (! alloc_real_seg(256, &para, &sel))
    fail("Couldn't allocate conventional memory"),

fp = (char far *) RC IP(sel, SRC_OFS),
_fstrncpy(fp, src); /* use prot-mode sel addr in program */
/* Generate real mode 21/60 */
memset(&r, 0, sizeof(r));
r.eax = 0x6000,
r.ebx = para; /* pass real mode para addr to DOS */
r.esi = SRC_OFS,
r.edi = DEST_OFS;
if (! dpmi_real_mode_intr(0x21, 0, 0, &r))
    fail("DPMI real mode interrupt failed"), //oops, should call free_real_seg
if (r.flags & 1) /* If carry set */
    _fstrncpy(dest, "invalid");
else
    _fstrncpy(dest, (char far *) RC IP(sel, DEST_OFS)),

/* copied to caller's dest, don't need conventional-memory
transfer buffer anymore */
if (! free_real_seg(sel))
    fail("Couldn't free conventional memory"),

return dest;
}

int PASCAL WinMain(HANDLE hInst, HANDLE hPrevInst,
LPSTR lpszCmdLine, int nCmdShow)
{
char buf[128], dest[128],
if (! (GetWinFlags() & WS_PRODE))
    fail("This program requires Windows Enhanced or Standard mode"),

if (! lpszCmdLine || ! *lpszCmdLine)
    fail("syntax: truname <pathname>"),
sprintf(buf, "TRJENAME %s", lpszCmdLine);
open_display(buf),
printf("%s", truname(lpszCmdLine, dest)),
show_display(),
return 0;
}

```

For example, if you add into the DOS command PROMPT (F4-F6) before starting Windows, then run `P TRJENAME F BAR` you'll produce the results `F100BAR`.

Note: in the protected mode version of truname.c, after allocating 256 bytes with `alloc_real_seg()` copies the contents of the conventional memory buffer, using the buffer's protected mode address. `truname()` then calls `INT 21h AH=00h` in real mode, passing in the buffer's real mode address. When `truname()` returns by way of the DPMI simulated real mode interrupt call,

truename copies DOS's reply to the caller's desk and frees up the transfer buffer. truename sends all this activity from its caller. It is almost as if DOS/MS-DOS implemented `FreeTransferBuffer` in the last pass.

### Windows and the PSP

There's something interesting about porting programs such as `TRUENAME` to Windows. After you get all the simulated real-mode interrupts and trap real calls and more correct, you run into something that isn't so obvious: there are things that you could do much more easily under DOS. You're going through a lot of extra work just to produce the same old results.

When porting old DOS utilities to Windows, don't do such a good job. You must consider how whatever you're working on may change under Windows. A good example is the classic MEM utility, which works under Eshay's the DOS Memory Control Block chain. Computer programs do extensive examples of such a utility. UDMEM, Natural, we could take a DMM as I'm not sure. Windows programs using `SetProcessWorkingSetSize` to create temporary protected-mode pointers in each MC. But what would a real example of that Windows memory management require different from the old plain vanilla DOS? A Windows version of UDMEM wouldn't even do anything about Windows and wouldn't need any new light on the MC's hardware.

One of the things the MEM utilities display are PSPs. If you do a hard port of UDMEM or a similar utility to Windows, you naturally should see PSPs for all the USRs, and for some Windows, and you see ones for core Windows execution, such as `WIN.COM`, `WIN386.EXE`, and `KERNEL386.dll`. But something is clearly missing. As noted earlier, every Windows task has an associated PSP. We'll then where are the PSPs. They can't be necessary, because the use of the `SetProcessWorkingSetSize` PSP. As we noted earlier, a Windows task opens files, so we can't see the PSP just as under real-mode DOS. Yet, as a single-minded Windows port of UDMEM would find the PSPs for Windows programs. Why? because where the Windows kernel allocates these PSPs—conventional memory, otherwise, they would be function of PSPs of some other program, or a large, or a system memory pool. Walking the MC chain is a good idea, but Windows isn't.

There is a separate source of Windows PSPs associated with the Windows task database chain and linked through task pointers. The `WINSPC` structure at offset 42 in the PSP. The Windows module responsible for running each of the old DOS programs, the Windows program, and previous structure of the offset 48 structure. PSPs according to DOS's basic string pointing to DOS look for a MEM entry in Windows's memory, these Windows entries only in the `WINSPC`.

Listing 3-21 shows a Windows PSP walk. `WINSPC`. This program does what you can using the normal MC to walk the chain of entries in Windows. The `Heap` library will be the address. `WINSPC` extracts pointers to the `IN` pointer from offset 60 in the `WINSPC` structure. `Undocumented Windows Chapter 5`. Finally, the program directly walks the Windows PSP chain, comparing back pointers stored at offset 42h.

### Listing 3-21 WINPSP.C Walks the Windows PSP Chain

```

/*
WINPSP.C
Andrew Schulman, April, 1993
From "Undocumented DOS", 2nd edition (Addison-Wesley, 1993)
Uses EasyWin:
gcc -W -Z winpss.c prot.c toolhelp.lib
*/

#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <dos.h>
#include "windows.h"
#include "prot.h"
#include "toolhelp.h"

```

```

#ifdef __cplusplus
extern "C" BOOL FAR PASCAL IsMinOldApTask(HANDLE hTask),
#else
extern BOOL FAR PASCAL IsMinOldApTask(HANDLE hTask),
#endif
#endif

#define MAP(ptr, bytes)      map_real(ptr), (bytes)
#define FREE_MAP(ptr)       free_mapped_linear(ptr)
#define GET_REAL(ptr)       get_real_addr(ptr)

#pragma pack(1)

typedef struct {
    BYTE type,
    WORD owner, /* PSP of the owner */
    WORD size,
    BYTE unused[3],
    BYTE name[8], /* in DOS 4+ */
} MCB,

WORD get_first_mcb(void)
{
    RMODE_CALL r,
    memset(&r, 0, sizeof(r)),
    r.eax = 0x5200,
    if (dpmi_rmode_intr(0x21, 0, 0, &r))
    {
        // Extract seg of first MCB from SysVars. Note that this is
        // at sysvars[-2]. You can't call map_real() on sysvars, and
        // then back up 2. You must map in sysvars-2 to begin with.
        WORD far *tmp = (WORD far *) MAP(REAL_PTR(r.es, (WORD)r.ebx-2), 2);
        MCB first_mcb = *tmp,
        FREE_MAP(tmp),
        return first_mcb,
    }
    else
        return 0,
}

WORD count_open_files(BYTE far *jft, WORD num_files)
{
    WORD open_files = 0,
    BYTE far *fp;
    int i,
    // count number of open files in a Job File Table (JFT)
    for (i=0, fp=jft; i<num_files; i++, fp++)
        if (*fp != 0xFF)
            open_files++,
    return open_files,
}

void display_psp WORD psp_seg, BYTE far *psp, MCB far *mcb)
{
    BYTE far *real_jft;
    BYTE far *prot_jft,
    WORD num_files;
    printf("204x:204x", psp_seg, mcb->size);
    if (_osmode >= 4)
    {
        char buf[9],
        _fmemcpy(buf, mcb->name, 8),
        buf[8] = '\0',
        if *buf)
            printf("\t-8s", buf);
    }
    real_jft = *((BYTE far *) far *) &psp[0x34];
    num_files = *((WORD far *) &psp[0x32]),
    prot_jft = (BYTE far *) MAP(real_jft, num_files);
}

```

```

printf(" \tZfp (Zu files, Zu open)\n",
    real_jft, num_files, /* print real mode addr */
    count_open_files(prot_jft, num_files)), /* use pmode addr */
FREE_MAP(prot_jft),
}

void fail(char *s) { printf("%s\n", s); exit(1); }

main()
{
    TASKENTRY te,
    BYTE far *maybe_psp,
    MCB far *mcb,
    WORD mcb_seg,
    WORD mapped,
    BOOL ok;

    printf( "-----\ndos apps:\n");
    // walk DOS MCB chain, looking for PSPs
    if (! (mcb_seg = get_first_mcb()))
        fail("Can't get MCB chain");
    for (;;)
    {
        mcb = (MCB far *) MAP(MK_FP(mcb_seg, 0), sizeof(MCB)),
        maybe_psp = (BYTE far *) MAP(MK_FP(mcb_seg + 1, 0), 512),
        if (*(WORD far *) maybe_psp == 0x20CD) // look like a PSP
        {
            if ((mcb_seg + 1) == mcb->owner) // regular DOS app PSP
                display_psp(mcb_seg + 1, maybe_psp, mcb);
        }
        FREE_MAP(maybe_psp),
        if (mcb->type == 'Z')
            break; // end of list
        mcb_seg = mcb_seg + mcb->size + 1; // walk list
        FREE_MAP(mcb),
    }
    FREE_MAP(mcb), // free last one
    mapped = get_mapped();
    if (mapped != 0)
        printf( ERROR! u mapped selectors remaining\n", mapped),
    printf("-----\nWindows apps:\n"),
    // now walk Windows task list, and extract PSPs
    te.duSize = sizeof(te);
    ok = TaskFirst(&te),
    while (ok)
    {
        BYTE far *tdb = (BYTE far *) MK_FP(te.hTask, 0),
        WORD prot_psp = *(WORD far *) tdb[0x60],
        WORD real_psp = GetSelectorBase(prot_psp) >> 4,
        if (*(WORD far *) MK_FP(prot_psp, 0) != 0x20CD) // really a PSP?
        {
            BYTE far *psp = (BYTE far *) MK_FP(prot_psp, 0),
            BYTE far *real_jft = *(BYTE far *) &psp[0x34],
            WORD num_files = *(WORD far *) &psp[0x32],
            BYTE far *prot_jft = (BYTE far *) MAP(real_jft, num_files),
            // Windows-specific fields in PSP
            WORD back_ptr = *(WORD far *) &psp[0x42],
            WORD flags = *(WORD far *) &psp[0x48],
            printf( "04X\t\t\t %s , real_psp, te.szModule),
            if (IsWinOldApTask(te.hTask) && !(flags & 1))
                fail("IsWinOldApTask flag weirdness"),
            putchar( (flags & 1) ? '*' : ' '); // IsWinOldApTask
            printf(" \tZfp (Zu files, Zu open)\n",

```

```

        real_jft, num_files, // print out real mode addr
        count open files(prot_jft,num_files)); // use pmode addr
    FREE_PSP(prot_jft);
}
ob = TaskNext(&ta);
}
printf("-----\nfollow PSP backlinks\n");
{
    WORD prot_psp, real_psp, next_psp,
    BYTE far *psp;
    prot_psp = GetCurrentPDB();
    while (prot_psp != 0)
    {
        real_psp = GetSelectorBase(prot_psp) >> 4,
        next_psp = ((WORD far *) %K_FP(prot_psp, 0x42)),
        printf("042 (%04X) -> %04X\n", prot_psp, real_psp, next_psp),
        prot_psp = next_psp,
    }
}
return 0;
}

```

Listing 3-7 shows sample WINPSP output. In addition to displaying the real mode paragraph addresses for each PSP and to name WINPSP shows the size of paragraphs of the MCB based on Windows NT's Srv. Since Windows PSPs don't have MCBs, the size field is blank. Any number here would be meaningless anyway since Windows programs can allocate gaps and gaps of extended memory. WINPSP also shows the address of the all important job table (JTB) associated with each PSP and displays the size of the JTB. Finally, WINPSP walks the JTB to determine the number of open files for each process.

**Figure 3-7: Sample Output from WINPSP**

```

DOS apps:
1180 0094 COMMAND      1180 0018 (20 files, 5 open)
1242 0102 DOSKEY       1242 0018 (20 files, 5 open)
1345 0200 WINICE       1345 0018 (20 files, 5 open)
1550 0055 win         1550 0018 (20 files, 5 open)
158F 0347 win386      158F 0018 (20 files, 5 open)
1900 8513 KRM386      278A 0000 (52 files, 18 open)
-----
Windows apps
257C 0018 (20 files, 5 open)
269C 0018 (20 files, 5 open)
4672 256A 0000 (30 files, 5 open)
2702 0018 (20 files, 5 open)
2734 0018 (20 files, 5 open)
-----
Follow PSP backlinks
11CF (2734) -> 1297
1297 (2702) -> 1247
1247 (269C) -> 1437
1437 (257C) -> 1817
1817 (4672) -> 00C7
00C7 (1900) -> 0000

```

WINPSP checks a Windows specific flag word at offset 40h in the PSP to determine when it has a DOS box; it marks these with an asterisk. Windows itself uses this field in the useful undocumented `IsWinOldApTask` function (see *Windows, and Windows*, Chapter 5). We use `IsWinOldApTask` as well as the ability to get at DOS boxes from Windows in Listing 3-28 below. WINPSP displays all

DOS boxes with the name "WINDOW" it would of course be more useful to display the name of whatever program is running in that DOS box, or at least the DOS box's window title.

### Peeking at DOS Boxes from a Windows Program

Speaking of DOS boxes, it would be nice if Windows programs could, in a easier time, getting at them. While it's a good thing that each virtual machine is a managed mode, has its own address space, a pointer such as 0x345678... in a DOS box has no real-world purpose (0x345678... in another DOS box or in a Windows program...), the other virtual machines that Windows programs communicate with, not Windows programs. Such communication would be extremely useful, given that Enhanced mode provides preemptive multitasking for some Windows programs. Windows programs are only cooperatively multitasking.

Fortunately, while each VM has its own segment of real address space, all the VMs are contained at a lower level in the low real address space. By mapping or linear addresses belonging to DOS boxes, a Windows program can peek or poke data in other VMs, likewise, protected mode DOS applications can map or linear addresses to access data in the Windows side of a VM.

For example, we might write a routine, `get_psp`, that is used by `WINSPY` to get the state of the program currently running in a DOS box. `get_psp` is suspended in `0x00000000` of a DOS box, so `get_psp` would need to get the `real_32` address of each DOS box, then, on each line of a file, `get_psp` would add an offset from the linear address for different DOS internal data structures and then create usable protected mode pointers to these structures, and the map linear... function from `PROT.C` (Listing 3-22). Listing 3-22 presents pseudocode for these steps.

### Listing 3-22: Pseudocode for Examining the Current PSP in Other VMs

```
sda = get_sda(),
for each DOS box
  high_lin = somehow get DOS box's high linear address, // see below
  // make linear address from SDA pointer
  sda_lin = PK_LIN(sda),
  // add together to get linear address of SDA in DOS box
  vm_sda_lin = high_lin + sda_lin;
  // make protected mode pointer to SDA in DOS box
  vm_sda = map_linear(vm_sda_lin, size of SDA),
  // extract current PSP from DOS box's SDA
  vm_curr_psp = vm_sda[0x10],
  free_mapped_linear(vm_sda), // done with SDA
  // get linear address of current PSP in DOS box
  #define SEG_TO_LIN(x) ((x) * 0x10)
  vm_curr_psp_lin = high_lin + SEG_TO_LIN(vm_curr_psp),
  // make protected mode pointer to current PSP in DOS box
  vm_psp = map_linear(vm_curr_psp_lin, size of PSP),
  // extract environment segment from DOS box's current PSP
  env_seg = vm_psp[0x2C],
  free_mapped_linear(vm_psp); // done with PSP
  // make protected mode pointer to current environment in DOS box
  vm_env_lin = high_lin + SEG_TO_LIN(env_seg),
  vm_env = map_linear(vm_env_lin, 0xFFFF);
  do usual stuff to walk env to find program name (see Chapter 7),
  display program name;
  free_mapped_linear(vm_env),
```

Let's go back that `get_psp` from a program running in one VM, either a Windows program or a DOS program, we want to be able to get at data in other VMs. For example, a Windows program might want to access DOS internal data structures, such as the SDA or `MSDOS.C` or `C:\PSY` in a DOS

boys. Because of the way that DOSMMGR works, many of these structures *must* be at the same real-mode segment offset address in each VM. As Chapter 1 noted, DOSMMGR assumes that even structures such as the CPU are not going to move. You can't alter this real-mode segment offset address of a linear address using the familiar *segment \* 16b + offset* rule. MK 11N (see Listing 3-17) to find the structure in a particular VM, you must add this linear address onto the VM's high linear address.

This segment offset address is the offset of the VM's address space within the entire 1-in-a-side mode address space. The key point is that this address is valid even when the VM is not running. For example, let's use a 2-MB linear address 12345h when a VM is running is always located at high lin = 12345h when the VM is not running. The current VM's memory has a base address of 0. We've been using some real addresses like that, but since we've always been accessing data in our own VM, without any other VMs around, we were able to ignore the high lin part. After all, if our program was running, this is VM was by definition the current VM.

So, every time you use any important high linear address for each VM (leaving aside for the moment how a program gets one of these addresses, note that debuggers such as WINICE display such addresses)

vm	vm handle	Status	High Addr	VM ID	Control Regs.
804c0000	00004000		81800000	00000002	80480f70
804b1000	0000f067		81400000	00000001	80010050

For example, in VMs, VM ID 1 is always the System VM (where all Windows applications run). Any other VM's high linear address, we can see, that the System VM's high linear address is the number 81800000, and that there is one DOS box, VM ID 2, which has the high linear address 81800000. Another example in this example, the DOS box's memory starts at 81800000b. A Windows program could peek at any place this memory is, always using linear = 0 at that address. You can now see why we need to map memory based on separate from map real. From here, commonly numbers you see in real mode addresses are not necessarily physical addresses (this machine doesn't have two gigabytes of memory).

Notice that WINICE displays *vm handle*, VM handles are unique identifiers for each VM. But a VM handle is more than that. The number happens to be the linear address of the VM's *control block*. The VM control block is a large data structure, only a few fields of which are documented (containing VM memory that the WINICE VM command-line options manage, including the VM's high linear address). I step with a window inside a VM CB structure that shows this and other fields. Given a VM handle, then one can get enough bits at address 0 to following code (I'm in for the "somehow get DOS box's high linear address" line in Listing 2-22).

```

DWORD vm_handle = somehow_get_vm_handle();
VM_ID far *vm_cb = map_l_near(vm_handle, sizeof(VM_CB));
DWORD high_lin = vm_cb->high_linear;
free_mapped_l_linear(vm_cb);

```

So, you'd like to see some more VM handles. This is the only difficult part. Fortunately, there is a way. VMM provides some mechanisms, documented in the SDK, that return VM handles (each one has a far VM handle, get near VM, Get Near VM Handle, follow far VM chain, and Get Sys VM Handle, get Next VM). Each of these functions returns a VM handle in the EBX register.

OK? Now, if only, there were some way we could call these functions from a Windows or DOS program. Unfortunately, VMM and VxD functions, *do not* appear in the WINDOWS.H header file. They do not yet seem to be part of the normal Windows programming repertoire.

Fortunately, you can still call these or any other VMM or VxD functions from a normal Windows or DOS program. Recall that normal programs *do not* or *can't* call VMM and VxDs all the time. VxDs



such as DOSMGR allocate callbacks and attach these to interrupts. Whenever a Windows program calls INT 21h, for example, it suddenly ends up running some 32-bit Real-Mode DOSMGR. Well, normal programs can't normally call into VxDs, too. Windows provides a service called Get Device Entry List Address structure (INT 11h, AX=1684h) which gives the ID number of a VxD returns a handle or pointer that a program can call to access services provided by the VxD.

This means that, for example, we could write a VxD to provide VM handle services to Windows and DOS programs. A program might call function 1 to get the current VM handle, function 2 to get the next VM handle and function 3 to get the system VM handle. The VxD would of course implement these functions by calling the VMM services there, acting as a surrogate for the Windows or DOS program. Clearly, VxDs aren't just for devices.

We can do even better than this, though. Why write a VxD just to get a three-function, 32-bit VM handle service? Why not get a different three functions. It is instead possible to implement a generic VxD which lets Windows and DOS programs access the function in VMM or a VxD. Here we use the generic VxD only to get at the three VMM handle functions provided by VMM. This barely shows off the capabilities of the generic VxD, which could be the subject of a whole book. For a taste of the rest of this, the generic VxD can do the explanation of how it works and a look at its source code, see "Can VxDs Factory an VMM Service? Easy Using Our Generic VxD," *Microsoft Systems Journal*, February 1993.

So how do we get the three services we get via handle pseudo-code above. It is provided by VMWALK.C, the VMWALK.H file, source code listings 3-23 and 3-24. VMWALK provides a service function which takes a function pointer, walkfunc. Using the generic VxD, VMWALK uses Get Sys VM Handle and Get Next VM Handle to enumerate all VM handles for a given virtual machine, vmwalk, gets the VM handle to form a VM CB, pointer copies it into a copy of the passed vmcb, and then writes it into the VM handle and a copy of a VM CB located at the high linear address, and can do what it wants.

### Listing 3-23: VMWALK.C

```
/*
VMWALK.C -- Enumerates VM handles, using generic VxD
Andrew Schulman, February 1993
From Undocumented DOS, 2nd edition (Addison-Wesley, 1993)
*/

#include <stdlib.h>
#include <string.h>
#include <dos.h>
#ifdef WINDOWS
#include "windows.h"
#endif
#include "prot.h"
#include "vxdcalls.h"
#include "vmwalk.h"

int vmwalk(WALKFUNC walkfunc)
{
    DWORD sys_vm = GetSysVMHandle();
    void far *vm_cb;
    VM_CB my_vm_cb,
    DWORD vm,
    int num_vm;
    for (vm=sys_vm, num_vm=0, num_vm++)
    {
        set_vm(0); // in case app fiddled with it
        vm_cb = map_linear(vm, sizeof(VM_CB),
            _frequency(&my_vm_cb, vm_cb, sizeof(VM_CB)),
            _free_mapped_linear(vm_cb),
            if (! (*walkfunc)(vm, &my_vm_cb))
```

```

        return num_vm;
    if ((vm = GetNextVRHandle(vm)) == sys_vm)
        break; // circular list
    }
    return num_vm;
}

DWORD GetSysVRHandle(void)
{
    VxDParams p;
    p.CallNum = Get_Sys_VR_Handle;
    return (VxDCall(&p)) ? p.OutEBX : 0;
}

DWORD GetCurVRHandle(void)
{
    VxDParams p;
    p.CallNum = Get_Cur_VR_Handle;
    return (VxDCall(&p)) ? p.OutEBX : 0;
}

DWORD GetNextVRHandle(DWORD vm)
{
    VxDParams p;
    p.CallNum = Get_Next_VR_Handle;
    p.InEBX = vm;
    return (VxDCall(&p)) ? p.OutEBX : 0;
}

// The next functions do nothing but add vm_cb->CB_High_Linear
// into a linear address, but they make it a lot easier to work with
// multiple VRs.
DWORD get_vm_base(DWORD vm)
{
    VR_CB far *vm_cb;
    DWORD base;
    if (vm == 0) return 0; // stay in current VR
    vm_cb = (VR_CB far *) map_linear(vm, sizeof(VR_CB));
    base = vm_cb->CB_High_Linear;
    free_mapped_linear(vm_cb);
    return base;
}

void set_vm(DWORD vm) // sets base address for map_linear
{
    set_base(get_vm_base(vm));
}

void far *map_real_vm(DWORD vm, void far *fp, WORD bytes)
{
    return map_linear(get_vm_base(vm) + RK_LINIFP, bytes);
}

void far *map_linear_vm(DWORD vm, DWORD lin_addr, WORD bytes)
{
    return map_linear(get_vm_base(vm) + lin_addr, bytes);
}

#ifdef TESTING
#include <stdio.h>
BOOL walkfunc(DWORD vm, VR_CB far *vm_cb)
{
    printf("0B1x VM #31x lin:0B1x crs:0B1x\n",
        vm,
        vm_cb->CB_VMID,
        vm_cb->CB_High_Linear,
        vm_cb->CB_Client_Pointer);
    return 1;
}
#endif

```

```

void fail(const char *s, ...) { printf("\n\n", s); exit(1); }
#endif OPNI_APP
#pragma argsused
int real_main(int argc, char *argv[]) { return 0; }
#pragma argsused
int pmode_main(int argc, char *argv[]) { return vmwalk(walkfunc); }
#else
int main() { return vmwalk(walkfunc); }
#endif
#endif /*TESTING*/

```

### Listing 3-24: VMWALK.H

```

/*
VMWALK.H -- See VMWALK.C
Andrew Schulman, February 1993
From "Undocumented DOS", 2nd edition (Addison-Wesley, 1993)
*/

// the first four fields are documented in DDK VMH.INC
typedef struct {
    DWORD CB_VR_Status,
    DWORD CB_High_Linear,
    DWORD CB_Client_Pointer,
    DWORD CB_VMID,
    BYTE other[0x100], // e.g., vm->next at VM+08h
    } VR_CB; // Virtual Machine Control Block

// Known VR_CB offsets in Windows 3.1
// 00h status
// 04h high linear
// 08h client pointer
// 0Ch VM ID
// 14h Get_Last_Updated_VR_Exec_File return
// 18h page tables
// 64h Get_Cur_PR_App_CB return
// 68h Get_Next_VM_Handle return
// BCh instance data struct
// 10Ch A20 wrap flag

// For more on the VR CB, see the two part article by
// Kelly Zytaruk in Dr. Dobbs's Journal, Dec 1993 and Jan 1994
#define NEXT_VM(vm_cb) *((DWORD far*)((BYTE far*)vm_cb+0x68))

typedef BOOL (*WALKFUNC) (DWORD vm, VR_CB far *);
int vmwalk(WALKFUNC walkfunc);
extern BOOL walkfunc(DWORD vm, VR_CB far *vm_cb); // app to provide
#endif OPNI_APP
// for Windows based VMWALK (without generic VxD, 3.1 Win apps only)
typedef BOOL (*WALK_FUNC) (HWND hwnd, char *title, DWORD vm,
    VR_CB far *vm_cb);
int vmwalk(WALK_FUNC walkfunc);
extern BOOL walkfunc(HWND hwnd, char *title, DWORD vm, VR_CB far *vm_cb);
#endif

DWORD GetSysVMhandle(void);
DWORD GetCurVMhandle(void);
DWORD GetNextVMhandle(DWORD vm);

void set_vm(DWORD vm); // sets base address for map_linear
void far *map_real_vm(DWORD vm, void far *fp, WORD bytes);
void far *map_linear_vm(DWORD vm, DWORD lin_addr, WORD bytes);

```

Given a VM handle, the `GetNextVMhandle` function returns the next VM in the list. `vmwalk` starts off its VM enumeration by calling `GetSysVMhandle`. The VM list is circular, so when `vmwalk` finds itself with the System VM handle (220) it is finished.

For how do `GetSysVMHandle` and `GetNextVMHandle` work in the first place? These functions are wrappers around calls to `VxDCall`. This function wrapper the code for it in a moment, is the C interface to the generic VxD. It expects to be passed a `VxDParams` structure, which contains all the information the generic VxD needs to make a VMM or VxD call on an application's behalf.

For example, the ODK documents `GetNextVMHandle` as taking a VM handle in the `EAX` register and returning the next VM handle, also in the `EAX` register. The `VxDParams` structure contains a page of the `EAX` register for both function input and output. It also contains a `CmdNo` field, which holds the function's magic number. `GetNextVMHandle` is defined in an `objfile.h` as `0x00003E`. These function numbers include the VxD ID number. VMM is considered to be VxD 1.

An `objfile` that wants to use the generic VxD must `#include VxDAPI.H`, and link with `VxDAPI.LIB` (see Listings 3.28 and 3.26). The generic VxD itself is the file `VxD386`, which must be installed with a file such as `device\windows2\chap3\vxd386` in the `[\386\obj]` section of `SYSTEM.INI`.

### Listing 3.25: VXD.CALLS.C

```
/*
VXD.CALLS.C -- C interface to generic VxD (VxD.386)
Andrew Schuman, February 1993
from undocumented DOS, 2nd edition (Addison-Wesley, 1993)
Abbreviated from version in Microsoft Systems Journal, February 1993
*/

#ifndef DPW1_APP
#include <windows.h>
#endif

#include <dos.h>
#include <vxdcalls.h>

static void far *API = 0;
static BOOL API_is_V86 = 1;
static char *requires_esq; /* This program requires VxD 386 */
extern void fail(const char *s, ...);

/* Get entry point for the VxD API using Int 2fh AX=1684h */
api_entry GetVxDAP1(MWORD vxd_id)
{
    _asm push di
    _asm mov ax, 1684h
    _asm mov bx, vxd_id
    _asm xor di, di
    _asm mov es, di
    qword int 2fh
    _asm mov ax, di
    _asm mov dx, es
    _asm pop di
    // returns in DI:AX
}

// This is only accurate with DPW1
BOOL IsProtMode(void)
{
    unsigned _ax;
    _asm mov ax, 1686h
    _asm int 2fh
    _asm mov _ax, ax
    return (!_ax);
}

/* check if generic VxD installed */
BOOL GenericVxDInstalled(void)
{

```

```

    return (GetVxDAPID(Generic_Dev_ID) != (api_entry) 0),
}

static void InitVxDAPID(void)
{
    if (! API)
    {
        if (! (API = GetVxDAPID(Generic_Dev_ID)))
            fail(requires_msg);
        API_is_V86 = (! IsProtMode());
    }
    else if (API_is_V86 && IsProtMode())
    {
        // API was last set in V86 mode, and we are now in protected mode
        API = GetVxDAPID(Generic_Dev_ID); // get PM API
        API_is_V86 = 0;
    }
}

BOOL VxDCall(VxDParams par *fp)
{
    InitVxDAPID();
    __asm push es
    __asm les bx, dword ptr fp
    __asm mov ax, VXD_VxDCall
    __asm call dword ptr [API]
    __asm pop es
    __asm jc Error
    return TRUE;
Error:
    return FALSE;
}

BOOL VxDPushCall(VxDPushParams par *fp)
{
    InitVxDAPID();
    __asm push es
    __asm les bx, dword ptr fp
    __asm mov ax, VXD_VxDPushCall
    __asm call dword ptr [API]
    __asm pop es
    __asm jc Error
    return TRUE;
Error:
    return FALSE;
}

```

### Listing 3-26 VXD.CALLS.H

```

/*
VXD.CALLS.H -- See VXD.CALLS.C
Andrew Schulman, February 1993
From "Undocumented DOS", 2nd edition (Addison-Wesley, 1993)
Abbreviated from version in Microsoft Systems Journal, February 1993
*/

#ifndef DPMI_APP
typedef int BOOL,
typedef unsigned short WORD,
typedef unsigned long DWORD,
#endif

#ifndef FALSE
#define FALSE 0
#define TRUE (! FALSE)
#endif

```

```

/* Structure used with VxDCall */
typedef struct {
    DWORD CallNum,
    DWORD Reserved1,
    DWORD InEAX, InEBX, InECX, InEDX, InEBP, InESI, InEDI,
    DWORD Reserved2, Reserved3,
    DWORD OutEAX, OutEBX, OutECX, OutEDX, OutEBP, OutESI, OutEDI,
    WORD OutFS, OutGS,
    DWORD OutEFLAGS,
} VxDParams,

/* Structure used with VxDPushCall */
typedef struct {
    DWORD CallNum,
    DWORD NumP,
    DWORD P[10],
    DWORD OutEAX, OutEDX,
} VxDPushParams,

#define CARRYFLAG 1
#define ZEROFLAG (1 << 6)
#define OVERFLOWFLAG (1 << 11)

/* VxD ID assigned by vxdid@microsoft.com */
#define Generic_Dev_ID 0x28c0
/* Functions supplied by VxD.386 */
#define VxD_Version 0
#define VxD_VxDCall 1
#define VxD_VxDPushCall 3

/*
W3MAP.H is a boring file generated from w3map -verbose
#define Get_VMM_Version 0x010000L
#define Get_Cur_VM_Handle 0x010001L
#define Test_Cur_VM_Handle 0x010002L
#define Get_Sys_VM_Handle 0x010003L
#define Test_Sys_VM_Handle 0x010004L

#define Call_When_Idle 0x01003AL
#define Get_Next_VM_Handle 0x01003BL
#define Set_Global_Time_Out 0x01003CL
.
etc.
*/
#include "w3map.h" /* available on disk */

/* function calls */
typedef void (far *api_entry)(void);
api_entry GetVxDAPID(DWORD vxd_id);
BOOL GenericVxDInstalled(void);
BOOL VxDCall(VxDParams far *fp);
BOOL VxDPushCall(VxDPushParams far *fp),

```

VxD.386 does little more than take the VxDParams structure passed in by a function such as `CallNextVMHandle` and pass the structure in EBP to the generic VxD entry point `VxDCall`, which in turn calls `GetVxDAPID` to get a device-specific entry point. It makes this necessary by calling `GetVxDAPID` + `GetVxDAPID` + 4 as a wrapper around `INT 2F` (see Table A-168-4). The generic VxD's ID is 28C0h. Microsoft assigns these VxD IDs.

With all the scaffolding in place, we are now ready to use the VMWALK facility to build a Windows program that examines data in DOS boxes. `FN MDRV.C` (shown in Listing 3-27) is getting tedious; it is a Windows application that displays the `CurrentDirectoryStructure.CDS` in all VMs. `FN MDRV` calls `w3mapk` to install a walktree, cleverly called `walktree`, which, for each VM, gets a pointer to the `CDS` and prints out all the drives and the current directory on each drive.

ENUMDRV uses the CURRDRIVE field from the SDA to indicate the current drive for each VM. As we know, there is a documented DOS function that returns the current drive by calling this function repeatedly for every drive for the caller that is for the current VM. It is possible to call DOS functions in other VMs, but it is considerably easier in this case—just group the other VM's SDA. So it turns out that the possibility back in Listing 3-4 about getting the CURRDRIVE value out of the SDA rather than calling the documented DOS function—which then just gets the value out of the SDA—maybe wasn't such a dumb idea after all.

### Listing 3-27: ENUMDRV.C Displays the CDS in All VMs

```

/*
ENUMDRV.C
Version of ENURDRV for multiple VMs in Windows Enhanced mode
Andrew Schulman, April 1993
from 'Undocumented DOS', 2nd edition (Addison-Wesley, 1993)

DPMI version:
bcc -2 -DDPMI_APP -c vxdcalls.c vmmwalk.c iswin.c
tlib dpmi_app+vxdcalls.obj+vmmwalk.obj+iswin.obj
bcc -2 -DDPMI_APP enumdrv.c dpmi_app.lib

Windows version (uses EasyWin)
bcc -2 -M -DWINDOWS enumdrv.c prot.c vxdcalls.c vmmwalk.c

WVRWALK version (uses EasyWin, no generic vsd)
bcc -M -DWINDOWS -DUSE_WVRWALK -2 enumdrv.c vmmwalk.c prot.c
*/

#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <dos.h>

#ifdef WINDOWS
#include "windows.h"
#else
#define DPMI_APP
#include "dpmish.h"
#endif

#include "prot.h"
#ifdef USE_WVRWALK
#include "vxdcalls.h"
#endif
#include "vmmwalk.h"

#define NETWORK (1 << 15)
#define PHYSICAL (1 << 14)
#define JOIN (1 << 13)
#define SUBST (1 << 12)
#define REDIR_NOT_NET (1 << 7) /* CDROM */

typedef unsigned char BYTE,
typedef unsigned short WORD,
typedef unsigned long DWORD,
typedef BYTE far *DPB, // provide actual DPB struct if needed

#pragma pack(1)
typedef struct {
    BYTE current_path[67], // current path (MAX_PATH != 67)
    WORD flags, // NETWORK, PHYSICAL, JOIN, SUBST
    DPB far *dpb, // pointer to Drive Parameter Block
    union {
        struct {
            WORD start_cluster, // root DDD0, never accessed ffff
            DWORD unknown,
        } LOCAL, // if (! (cds[drive] flags & NETWORK))
        struct {

```

```

        DWORD redirect_record_ptr;
        WORD parameter,
        ) NET; // if (cds[drive].flags & NETWORK)
    } wj
    WORD backslash_offset, // offset in current_path of '\
    // DOS4 fields for IFS
    // 7 extra bytes...
    ) CBS; // Current Directory Structure

CDS far *curdir(unsigned drive);
void set_vm(DWORD vm),
#define ENHANCED MODE 3
// see [SWIN [
extern int is_win(int *pms), int *pms, int *pms);
void fail(const char *s,
{
    puts(s),
#ifdef DPMI_APP
    _dos_ext(1),
#else
    exit(1),
#endif
}
static void far *sysvars_real = 0;
static void far *sdm_real = 0;
static int curdir_size = 0;
#ifdef DPMI_APP
// going to do initialization from real mode (real_main())
#define CALL_DOS(x,y,z) int86s(0x21, x, y, z)
#else
#define CALL_DOS(x,y,z) real_int86s(0x21, x, y, z)
#endif
void far *get_sysvars(void)
{
    union REGS r,
    struct SREGS s,
    memset(&r, 0, sizeof(r)),
    memset(&s, 0, sizeof(s)),
    r.h.ah = 0x32,
    CALL_DOS(&r, &r, &s),
    return HK_FP(PS es, r.b);
}
void far *get_sdm(void)
{
    union REGS r,
    struct SREGS s,
    memset(&r, 0, sizeof(r)),
    memset(&s, 0, sizeof(s)),
    r.x.ax = 0x000,
    CALL_DOS(&r, &r, &s),
    return (r.x.cflag) ? (void far *) 0 : HK_FP(s.ds, r.b);
}
void do_init(void)
{
#ifdef USE_VIRTUAL
    if (! GenericVxDInstalled())
        fail("This program requires the generic VxD (VED.386)\n"
            "Put device=vxd.386 in [386Enh] section of SYSTEM.INI");
#endif
    // assumes SysVars at same address in all VMs
    // but CDS not necessarily at same address

```



```

if (! (sysvars_real = get_sysvars()))
    fail("21/52 failed: can't get SysVars'\n");
// assumes SDA at same address in all VMs
if (! (sda_real = get_sda()))
    fail("21/5006 failed: can't get SDA'\n");
// problem! Microsoft C QuickWin defines _osmajor, _osminor
// with the Windows version, not the DOS version!!
_asm mov ax, 3000h
_asm int 21h
_asm mov byte ptr _osmajor, al
currdir_size = (_osmajor >= 4) ? 0x58 : 0x51,
)
#endif WINDOWS
int main(int argc, char *argv[])
{
    if (! (GetWinFlags() & WF_PROTE))
        fail("This program requires Windows 3.x Enhanced mode");
    do_init();
#ifdef USE_VVMWALK
    return vvmwalk(walkfunc);
#else
    return vmwalk(walkfunc);
#endif
}
int real_main(int argc, char *argv[])
{
    int maj, min, mode;
    puts("EHUNDRV - Enumerate OS in all Windows VMs,
    puts(" from 'Undocumented DOS', 2nd edition (Addison-Wesley, 1993)\n");
    if (! (is_win(&maj, &min, &mode) || (maj < 3) ||
        (mode != ENHANCED_MODE)))
        fail("This program requires Windows 3.x Enhanced mode");
    do_init(); // call in real mode
    return 0; // okay to switch into protected mode
}
int pmode_main(int argc, char *argv[])
{
    vmwalk(walkfunc);
    return 0;
}
#endif
#ifdef USE_VVMWALK
BOOL walkfunc(HWND hwnd, char *title, DWORD vm, VR_CB far *vm_cb)
#else
BOOL walkfunc(DWORD vm, VR_CB far *vm_cb)
#endif
{
    CBS far *cbs, far *dir;
    BYTE far *sysvars,
    BYTE far *cbs_real,
    BYTE far *sda,
    DWORD total, notmapped,
    int lastdrv, curdrv,
    int i,
    static DWORD sys_vm = 0;
    static DWORD cur_vm = 0;
#ifdef USE_VVMWALK
    if (! sys_vm) sys_vm = GetSysVMMHandle();
    if (! cur_vm) cur_vm = GetCurVMMHandle();

```

```

endl;
#define CDS_OFS      0x16
#define LASTDRV_OFS 0x21

// Set the base address (could also have used map_linear_vm
// or map_real_vm. This one call to set_base will affect all
// subsequent calls to map_linear.
set_base(vm_cb->CB_High_Linear);
sysvars = (BYTE far *) map_real(sysvars_real, LASTDRV_OFS+1),
// Note does not assume that CDS at same address in all VMs'
// But moving a CDS breaks in Windows anyway, because DOSMGR assumes
// that CDS never moves after Windows initialization. *Sigh*
cds_real = (BYTE far * far *) &sysvars[CDS_OFS],
lastdrv = sysvars[LASTDRV_OFS],
free_mapped_linear(sysvars);

printf("\n%s (%VM %lu) -- CDS at %08lx\n",
      (vm == sys_vm) ? "System VM"
      (vm == cur_vm) ? "Current VM"
      /* default */ "DOS Box ",
      vm_cb->CB_VMID,
      vm_cb->CB_High_Linear + PK_LIN(cds_real)),

#define CURRDRV_OFS 0x16

// get current drive for this VM from SDA:0016
// for System VM, will depend on current task
// each Windows app has its own current drive/directory
// a ready did set base vm_cb->CB_High_Linear),
sda = (BYTE far *) map_real(sda_real, CURRDRV_OFS+1),
currdrv = sda[CURRDRV_OFS],
free_mapped_linear(sda),

// map 'n cds_real for this VM, print out, free it
for (CDS far *) map_real(cds_real, lastdrv * currdrv_size),
for (i = 0; i < lastdrv; i++)
{
    dir = (CDS far *) ((BYTE far *) cds) + (i * currdrv_size),
    if (dir->flags)
    {
        // highlight current drive for this VM
        printf("%i == currdrv) ? "map" : " " );
        printf("%c\t\t %0Fs", 'A' + i, dir->current_path);
        if (dir->flags & REDIR_NOT_NET) printf(" REDIR_NOT_NET ");
        else if (dir->flags & NETWORK) printf(" NETWORK ");
        if (dir->flags & JOIN) printf(" JOIN ");
        if (dir->flags & SUBST) printf(" SUBST");
        printf("\n");
    }
}
free_mapped_linear(cds),
set_base(0); // restore
return TRUE;
}
}

```

Figure 3-8 shows output from ENUMDRV. By looking at all the VMs, this program provides a telling diagnosis: two main Windows main tasks, multiple current drives and directories by instantiating the CDS. The CDSs sit at the same memory offset address with each VM, though ENUMDRV does not in fact depend on this. But it says that each Windows application is its own contrasting this with, for example, the SF1, which is not instantiated, ask from the mix to the PC VMs and GrowSF1 segments, discussed earlier.

**Figure 3-8: ENUMDRV Shows the CDS in Each Windows VM**

```

System VM (VM 1) -- CDS at 814100B0
A A:\

```



```

// from Undocumented Windows, p. 111
BOOL IsDosBox(HWND hwnd)
{
    if (! IsWinOldAppTask) // one-time init
        * ( ( IsWinOldAppTask = (BOOL) (FAR PASCAL *) (HANDLE)
            GetProcAddress(GetModuleHandle("KERNEL"), "ISWINOLDAPTASK"))
            ? "Can't find KERNEL IsWinOldAppTask";
    return (* IsWinOldAppTask) (GetWindowTask(hwnd));
}

typedef BOOL (*WALKFUNC)(HWND hwnd, char *title, DWORD va, VR_CB far *va_cb);
int vswalk(WALKFUNC walkfunc)
{
    char buf(0x50),
    VR_CB dummy,
    HWND hwnd,
    int num_va;

    if (((WORD) GetVersion()) != 0x0403)
        return 0; // only works in Windows 3.1
    if (! (GetWinFlags() & WF_ENHANCED))
        return 0; // only works in Enhanced mode

    hwnd = GetActiveWindow(),
    GetWindowText(hwnd, buf, 0x50),
    // Can't get handle for System VR, but that is VR the Windows
    // app is running in! So make a dummy VR_CB and pass that
    // (actually, COULD get System VR by walking VR next ptr):
    memset(&dummy, 0, sizeof(dummy));
    dummy.cb.High_Linear = 0; // current VR!
    dummy.cb.VRID = 1; // System VR is always VR 1
    if (! (*walkfunc)(hwnd, buf, 0, &dummy))
        return 0;
    num_va = 1;
    hwnd = GetWindow(hwnd, GW_HWNDFIRST),
    while (hwnd)
    {
        if (IsDosBox(hwnd) && GetWindowText(hwnd, buf, 0x50))
        {
            DWORD va = VR_FROM_Hwnd(hwnd),
            VR_CB far *va_cb (VR_CB far *)
                map_linear(va, sizeof(VR_CB)),
            VR_CB far va_cb,
            *memcpy(&va_cb, va_cb, sizeof(VR_CB)),
            *free_mapped_linear(va_cb);
            if (! (*walkfunc)(hwnd, buf, va, &va_cb))
                return num_va;
            set_base(0); // in case app changed base
        }
        hwnd = GetWindow(hwnd, GW_HWNDNEXT),
    }
    return num_va;
}

#ifdef TESTING
#include <stdio.h>
BOOL vswalkfunc(HWND hwnd, char *title, DWORD va, VR_CB far *va_cb)
{
    printf(" 0x%04x: 0x%08lx: 0x%08lx\n", hwnd, va, title);
    printf(" VR #1:u st=0x%08lx cin=0x%08lx crs=0x%08lx\n",
        va_cb->CE_MID,
        va_cb->CB_VR_Status,
        va_cb->CB_High_Linear,
        va_cb->CB_Client_Pointer);
}

```

```


    return t;
}
main()
{
    return wvwalk(wwalkfunc);
}
endif

```

Aside from its extreme version dependence, there is a problem with WYMWALK. It does not get a handle for the System VM. However, this is not a major problem. The technique only works for Windows applications, not Windows application instances, the System VM. But, remember, when a Windows application sees a WYMWALK version of ENUMDRV, it runs as if on a System VM. The current VM to the program does not need to be high in the address space. It is accessible by CDN or whatever with a base offset, just as we did a long when we wrote the code to know about VM handles. High base addresses and other details WYMWALK works as a virtual VM. Call for the System VM. For the more, by using the next pointer stored at 006c:0080 to each VM. Call for the VM linked to it. It is possible to create the System VM which is VMID 01.

Actually, you could even see how to use a VM to ROM (HWND) technique to connect to with the generic VxD. In Fig. 3-28, a would like to have some more information for each VM, such as its window title. Meanwhile, WYMWALK version since it starts with a window title, is able to display window titles (see Fig. 3-9). You could connect the process context to produce a single, more useful display. After all, most users don't think of VMs or even of DOS boxes, they think of the applications they are run on.

Figure 3-9: WYMWALK Version of ENUMDRV



Note finally that ENUMDRV can also be built as a DPMI\_APP. There is a twist here, however. We earlier built several utilities for protected DPMI, and these could actually run not only in Windows but in a DPMI mode such as 386MAX 6.0 and 6gh. With ENUMDRV, on the other hand, we have a DOS program that sees a WYMWALK and VXD Allen, really does require Windows Enhanced mode. ENUMDRV sees the function is win\_ to ensure it is running in protected mode. Listing 3-29 shows ISWIN.C.

## Listing 3 29: ISWINLC: DOS Code to Check for Windows

```

/*
ISWIN.C
Detecting Windows mode, version from DOS
Andrew Schulman, February 1993
bcc 0STANDALONE iswin.c
*/

#include <stdlib.h>
#include <stdio.h>
#include <dos.h>

#define REAL_MODE      1
#define STANDARD_MODE 2
#define ENHANCED_MODE 3
#define SYSTEM_VM     1

int detect_switcher(void)
{
    int retval = 1;
    _asm push di
    _asm push es
    _asm xor bx, bx
    _asm mov di, bx
    _asm mov es, bx
    _asm mov ax, 4002h
    _asm int 2fh
    _asm mov cx, es
    _asm or cx, di
    _asm je no_switcher
done:
    _asm pop es
    _asm pop di
    return retval;
no_switcher:
    retval = 0;
    goto done;
}

// a lot more complicated than you would have thought!
int is_win(int *maj, int *min, int *mode)
{
    unsigned short retval;
    int maj=0, min=0, mode=0;

    /* make sure someone, anyone has INT 2fh */
    if (_dos_getvect(0x2F) == 0)
        return 0;

    /* call 2F/160A to see if Windows 3.1+ */
    _asm mov ax, 160Ah
    _asm int 2fh
    _asm mov retval, ax
    if (retval == 0) /* AX=0 if Windows running */
    {
        _asm mov mode, cx /* CX=2 means Std; CX=3 means Enh */
        _asm mov byte ptr maj, bh /* BX = major/minor (e.g., 030Ah) */
        _asm mov byte ptr min, bl
        *maj = maj;
        *min = min;
        *mode = mode;
        return 1;
    }

    /* call 2F/1600 to see if Windows 3.0 Enhanced mode or Windows/386 */
    _asm mov ax, 1600h
    _asm int 2fh
    _asm mov byte ptr maj, al

```

```

_asm mov byte ptr min, ah
if ((maj == 1) || (maj == 0xFF)) /* Windows/386 2.x is running */
{
    *pmsj = 2; /* Windows/386 2.x */
    *pmin = 1; /* don't know; assume 2.1? */
    *pmode = ENHANCED_MODE; /* Windows/386 sort of like Enhanced */
    return 1;
}
else if ((!(maj==0) || (maj==0x80))) /* AL=0 or 80h if no Windows
/* must be Windows 3.0 Enhanced mode */
{
    *pmsj = maj,
    *pmin = min,
    *pmode = ENHANCED_MODE;
    return 1;
}

/* call 2F/4680 to see if Windows 3.0 Standard or Real mode, but,
this could be a "3.0 derivative" such as DOSSHELL task switcher! */
_asm mov ax, 4680h
_asm int 2fh
_asm mov retval, ax
if (retval == 0) /* AX=0 if 2F/4680 handled */
{
    /* make sure it isn't DOSSHELL task switcher */
    if (detect_switcher())
        return 0;
    *pmsj = 3;
    *pmin = 0,
    /* either have Windows Standard mode or Real mode; to
distinguish, have to do fake Windows broadcasts with
2F/1605 'xak' We'll avoid that here by assuming
3.0 Standard mode if you really want to distinguish
3.0 Standard mode and Real mode, see _MSJ_, March 1991,
p. 113, and MS KB articles Q75943 and Q75338 */
    *pmode = STANDARD_MODE;
    return 1;
}

/* still here — must not be running Windows */
return 0;
}

#ifdef STANDALONE
main()
{
    int maj, min, mode=0,
    if (! is_win(&maj, &min, &mode))
        printf("Windows is not running\n");
    else if ((maj == 2))
        printf("Running windows/386 2.x\n");
    else
        printf("Running Windows 3u (02u or higher) %s mode\n",
            maj,
            min,
            (mode == REAL_MODE) ? "Real" :
            (mode == STANDARD_MODE) ? "Standard" :
            (mode == ENHANCED_MODE) ? "Enhanced" :
            /* don't know */ "???");
    if (mode == ENHANCED_MODE)
    {
        unsigned short vax;
        /* call 2F/1683 to see if DOS app running in System VR; if so,
this must be some hacked version of Windows like MSDPM1,
or we must be running inside WINSTART.BAT */
        _asm mov ax, 1683h
        _asm int 2fh
    }
}

```

```

asm mov vm, bx
if (vm == $SYSTEM_VIR)
    printf("Running DOS app in System VR
        "Must be WINSTART BAT or hacked-Windows!\n");
else
    printf("VR #Zulu", vm);
}
/* could also call 21/160C to check for Windows in ROM */
return mode, /* can be tested with ERRORLEVEL */
}
#endif

```

Something as seemingly simple as checking for Windows turns out to be so complicated that it requires 150 lines of code!

## Undocumented DOS and DesqView

by Rolf Brown

**With all this talk about Windows**, what about other multitasking windowing environments for DOS, such as Quarterdeck's DESQview?

Although no special preparations are required to access undocumented DOS functions or data from a program running under the DESQview multitasker, task switching should be disabled when modifying DOS data structures through any method other than INT 21h calls.

DESQview analyzes access to DOS via INT 21h, thus automatically avoiding problems with DOS data structure modifications. If a program directly modifies DOS internal data, however, a task switch could occur when the data is being modified. Ordinarily, one could just disable interrupts while making the modifications, but this will not work on 386 and higher processors. QEMM (used in DESQview 386) and other memory managers virtualize the interrupt flag with the result that DESQview is able to task switch even while interrupts appear to be disabled. Therefore, task switching must be explicitly disabled by asserting a critical section using TopView API calls. While IBM's TopView has gone to the land of the forgotten, here's your chance to purchase TopView for \$4.95 at a store called Weird Stuff in Sunnyvale. (A TopView SIM then costs \$7.95); its API still lives on in DESQview, as well as in Sunny Hill Software's TaskView/OverView.

### Listing 3-30: DESQVIEW.C (Actually, More Like TopView)

```

static int TopView = 0;
void check_TopView(void)
{
    union REGS regs;
    regs.x.ax = 0x1022; /* get TopView version */
    regs.x.bx = 0;
    int86(0x15, &regs, &regs);
    TopView = regs.x.bx; /* nonzero if TopView or compatible */
}
void TopView beginCrit(void)
{
    union REGS regs;
    if (TopView)
    {
        regs.x.ax = 0x101B; /* start critical section */
        int86(0x15, &regs, &regs);
    }
}

```



```

void TopView_endcrit(void)
{
    union REGS regs ;
    if (TopView)
    {
        regs.x.ax = 0x1010 ;           /* end critical section */
        int86(0x15,&regs,&regs) ;
    }
}

```

These three functions are used in the following manner: call `check_TopView()` once near the beginning of the program to initialize the `TopView` variable. Then, each time a DOS data structure is to be modified, surround the code performing the modification by calls to `TopView_begincrit()` and `TopView_endcrit()`, which disable and enable task switching, respectively.

This chapter's discussion of Windows focused entirely on *protection* at DOS internal data structures, so it's a shame that critical sections never came up. However, Raj's point about *modeling* structures in DPMI view actually applies with equal force to Windows, at least when a DOS box's Background execution is used. Programs that modify DOS internal data should use the Background Section IN 210-AN-1681a and real-time Section IN 211-AN-1682a, which the Windows DDK documents. Do not forget that a useful Microsoft Knowledgebase article "Using the Interrupt 21 Critical Section Services" (Q78181) notes that, despite their names, these functions "do not prevent a task switch... occurring, and impact task switching only in a limited way." Success! It's a good idea, though, to guarantee that a response could be given whether ripped and switched, was torn, connected, or put the code in a VxD.

## A Brief Introduction to VxD Programming

Speaking of VxDs, we've seen so many VxD code fragments, both here and in Chapter 4, that we really need to look at them as complete, substantial VxDs.

UNDOSMGR.ASM (using VxDs) is the 32-bit assembly language source code for UNDOSMGR.386, a small, simple, but useful program that provides support for protected mode in a single task environment. For more details, see IN 211-AN-51006a (see VxDs: Externals). As there, we saw how a program running under Windows can use DPMI to call this function from protected mode. Because this program must explicitly use DPMI to access the hardware, this access supports a of course, non-transparent.

However, if UNDOSMGR.386 is installed without one, such as `device=undoc2setup3-undocmgr.386` (i.e., 386) in a system of `SYSTEM.INI` (subfile IN 211-AN-51006a), supported in protected mode, and nonprotected mode programs no longer need to access the DPMI. They can just call the hardware directly from protected mode, and the UNDOSMGR.VxD takes care of everything. With UNDOSMGR installed, a straight-protected mode port of CURRDIR.C (Listing 3) will work properly, and handle I/O of longer arguments (i.e., messages) from the debug version of DOSMGR.

### Listing 3 31: UNDOSMGR.ASM

```

; UNDOCMGR.ASM
; Sample vxd that provides one undocumented DOS function (21/5006)
; In protected mode
; Andrew Schulman, April 1993
; from Undocumented DOS, 2nd edition (Addison-Wesley, 1993)
; 386p

INCLUDE VMM.INC
INCLUDE v86MMG.INC

```

```

.....
, VIRTUAL DEVICE DECLARATION
.....
Declare_Virtual_Device UNDOSMGR, 1, 0, \
Control_Proc, \
Undefined_Device_ID, \
Undefined_Init_Order, , , \
,
, ONLY DATA SEGMENT
,
.....
VxD DATA_SEG
Prev_Int21Pmode dd 0
Prev_Int21Pmode_Seg du 0
VxD DATA ENDS
,
, MOVEABLE CODE SEGMENT
,
VxD_CODE_SEG
BeginProc Int21Pmode
mov ebx, [ebp.Client_EAX]
cmp ebx, 5D00h
je short Do_GetSDA

Default:
movzx ecx, Prev_Int21Pmode_Seg ;, "NOT" mov cx"
mov edx, Prev_Int21Pmode
VMMcall $, m, ate_fer_jmp
jmp short Fini

Do_GetSDA:
IFDEF BYPASS_VSBMMGR
, sample code to show how this looks with raw VMM calls
VMMcall Simulate_Instr
VMMcall Begin_Nest_VB6_Exec
mov esi, 71h
VMMcall Exec_Int ; reflect INT 21h to VB6 mode
VMMcall Resume_Exec ; do it now?
movzx ebx, [ebp.Client_DS] ; get VB6 DS inside nest
VMMcall End_Nest_Exec
shl ebx, 4 ; make linear address
mov ecx, 0FFFFh ; 64k segment
VMMcall Map_Lin_To_VB6_Addr ; create permanent selector
mov [ebp.Client_DS], cx ; return selector to caller in DS
ELSE
VMMcall Simulate_Instr
mov edx, OFFSET$ GetSDA_API
VxDcall VSBMMGR_Xlat_API
ENDIF
Fini:
clic
ret

GetSDA_API:
xlat_API_Return_Ptr DS, $!
xlat_API_Exec_Int 21h
EndProc Int21Pmode
VxD_CODE_ENDS
,
, LOCKED CODE SEGMENT
,
VxD LOCKED_CODE_SEG
BeginProc Control_Proc

```

```

Control_Dispatch Device_Init, Do_Device_Init
clic
ret
EndProc Control_Proc
VxD_LOCKED_CODE_ENDS
;
; ***** PM INIT CODE SEGMENT *****
;
VxD_ICODE_SEG
BeginProc Do_Device_Init
; get previous pmode INT 21h handler, so can chain
mov eax, 21h
VMCall Get_PM_Int_Vector
mov Prev_Int21Pmode, edx ; offset
mov Prev_Int21Pmode_Seg, cx ; segment
; turn int21 handler into pmode callback
xor edx, edx
mov esi, OFFSET$2 Int21Pmode
VMCall Allocate_PM_Call_Back
jc $short Done
; segment:offset callback address in EAX
; install pmode int21 handler
mov ecx, eax
movzx edx, ax
shr ecx, 10h
mov eax, 21h ; eax = inno
VMCall Set_PM_Int_Vector
Done:
ret
EndProc Do_Device_Init
VxD_ICODE_ENDS
;
; ***** REAL MODE INITIALIZATION *****
;
VxD_REAL_INIT_SEG
; just a stub real_init procedure
real_init proc near
xor ax, ax
xor bx, bx
xor si, si
xor edx, edx
ret
real_init endp
VxD_REAL_INIT_ENDS
END

```

VxDsrc 32-bit linker is available. If you wish to use the standard Microsoft assembler and linker currently can't be used to merge UNDISMGR.ASM into UNDISMGR.386 you will need Microsoft's MASM5 LINK 386 and ADDRESSOP utilities and Microsoft's all important VMMEMH.C header file. Included with the Windows SDK and with the "VxD Inc." package available on the MSDN CD-ROM. MK\_VXD is a batch batch file for making simple VxDs.

```

ren MK_VXD.BAT
set incude=\\ddk\inc\ude
masm5 -p -u? X2 X1,
link386 X1,X1 386,.,,? def
addr? X1,386

```



```

#include <stdlib.h>
#include <stdio.h>
#include <dos.h>
#include <time.h>
#include "dpwish.h"

void dospeed(unsigned long iter)
{
    time_t t1, t2;
    unsigned long i;
    time(&t1);
    for (i=0; i<iter; i++)
    {
        _asm mov ah, 2Ah          /* get date */
        _asm int 21h
        _asm mov ah, 2Ch          /* get time */
        _asm int 21h
        _asm mov ah, 51h          /* get PSP */
        _asm int 21h
        _asm mov ax, 3000h        /* get version */
        _asm int 21h
    }
    time(&t2);
    printf("21u calls in 21u seconds\n", iter * 4, t2 - t1);
}

void fail(const char *s, ...) { puts(s); _dos_exit(1); }
int real_main(int argc, char *argv[])
{
    printf("real mode: ");
    dospeed((argc < 2) * 10000 + atol(argv[1]));
    return 0;
}

int prot_main(int argc, char *argv[])
{
    printf("prot mode: ");
    dospeed((argc < 2) * 10000 + atol(argv[1]));
    return 0;
}

```

On a standard 80386SX 20-MHz 486 DMMSPED turned in the following times for 10 000 iterations of the four calls: INT 21h-AH 2Ah-AH 2Ch-AH 30h and AX with 3000h from our earlier discussion that these BIOS functions are typically called very frequently.

MS-DOS 5 (no memory manager)	5 seconds
MS-DOS 5 (386MAX or EMM386)	8 sec
MS-DOS 5 (Windows Enhanced mode DOS box)	13 sec
386MAX DPWMI client	23 sec
Windows Enhanced mode DPWMI client	26 sec
W Enh mode, with UNDOSMGR 386 installed	29 sec
Novell DOS 7 with KRN.386 SYS	20+ sec
Novell DOS 7, in DPWMI	50+ sec

As with a 4-m-hastily-constructed benchmark, one can conclude from these figures whatever one wants. However, I do see some wariness that there is a small but noticeable overhead for installing VxDs that hook frequently used software interrupts and that more importantly, there is a large and very noticeable overhead for running V86 mode under Windows Enhanced mode and in every regard and more noticeable overhead for making I/O calls from protected mode. As a fitting conclusion to this chapter on putting DOS programs to protected mode, perhaps DMMSPED is trying to tell us that as long as Windows has to call down to real mode DOS, the best way to make a lot of DOS calls from protected mode is not to make them at all.



## Other DOSs: From DR DOS and NetWare to MVDMS in OS/2 and Windows NT

for better or worse. In MS-DOS 5.0's programer-to-programer handshake, a worldwide computer standard (which by some was non-programmer languages such as C or C++) standards, there is no committee to oversee the DOS standard; it is a proprietary standard controlled by Microsoft. But many other companies have implemented the DOS interface. DOS is not a desktop-only system; many operating systems feature MS-DOS, including UNIX OS/2 and Microsoft's own Windows NT.

The same thing is starting to happen with Windows: installable components are now under way, and UNIX manufacturers such as the Windows Application Library Interface (WALI) from Sun Microsystems, NetWare's *doswin*, and the early *doswin* Project 1 technologies. In fact, Sun is promoting its desktops. Labels Windows-compatible (PWI) found in Windows 4.11 are the public domain. *DR-DOS for DOS* (Veritas, 1988, price \$4.95) and *NetWare* (Novell) are the first. IBM QuantumLink WordPerfect and other companies have also had an interest in how to implement PWI (but one could argue that DOS and Windows are DOS's after all, given the amount of software required to design something like ANSI standards committees).

Even when DOS is not the best choice, companies need the DOS interface. DOS is not always Microsoft's MS-DOS. In addition to several of the companies such as Veritas and Central Software, which produce clones of workalikes of DOS for the computer systems market, there is Novell's *Price-Up* (see *see NetWare*, Feb 1991), Novell's *price-up*, Digital Research's (DRI) original *dr* file, Co. Manufacturing systems that create MS-DOS-like (but not DOS) and users of a computer to MS-DOS-like (DR-DOS, DR-DOS has seen several Novell DOS clones that come out prior to Novell's NetWare, the only operating system market Novell DOS has produced since the early 1980s) version of Microsoft's MS-DOS.

What if it is a matter of how DOS is implemented, rather than how you know that you can't just run a Microsoft's MS-DOS program? *DR-DOS* is a good example of this. It provides a fair self-imposed compatibility with a DOS implementation, almost as popular as a publication was in. Looking just at the INTEL architecture, you have to implement the *dos* system function 52h. And this was a major engineering effort, just some of the *dos* system. List of I/O structure. If you wanted to run any new software, you would have to implement the DOS-like of the network reduced overhead (INTEL-M144, see Chapter 8). And if you want to run Windows, and therefore run WinWord and Excel, you would have to copy Swapfile, Data Area and location \$D00h (plus a few other stuff that would be implementation details, but which are part of the *de facto* DOS specification).

In implementing a version of DOS, a company might have some good ideas about how to implement different data structures. Tough luck. You had better have an MS-DOS-style Current Directory Structure for the above-mentioned tasks, a System File Table for open files, Memory Control Blocks, and so on. At least you need copies of them, enough to keep different applications happy, and then some mechanism for transferring state from these dummy DOS structures to whatever internal structures your system actually uses.

Even worse if you were creating a DOS-workalike, you would need to keep a number of support bytes for a variety of critical locations in your DOS's data segment. For example, later in this chapter we find that some Microsoft VC compilers depend on the presence of a flag at offset 4 to the DOS data segment and expect the current PDB to be stored in the same segment at either offset 210h or 330h (see listing 4-4).

All of the interplay between DOS applications and DOS is very wide. The DOS specification itself is really the application owners themselves. Any self-respecting DOS-workalike must, by definition, support whatever the popular PC applications do.

As the problem of implementing an operating system into specific hardware requirements is not unique to DOS, a vendor who is working in the competitive business, trying to get applications from one environment to another, will probably know how you go about such a situation. You take the real applications, find their variations in structure, and you try to run them. When the applications don't or when something goes wrong, you make some change or addition to your operating system.

The specification of a workalike is collected from the applications. You see what the applications expect of their environment, and you provide it.

Let's start our work of essays on software engineering. *The Mirror of Man-Month*, Fred Brooks has a good discussion of compatibility with *de facto* standards.

An implementation exercise is a form of definition. When the first compatible computer was built, it was to excel the technology used. The new machine was to match an existing machine. The standard was copy or some poeas. Ask the machine's. A test program would be developed to determine the behavior, and the new machine would be built to match.

Using a compatible form from a different source has advantages. All questions can be set back to the original development. Deficiencies are worked out, answers are quick. Answers are always positive, since you're adding to a well-established definition. Opposed to this, one has a formidable set of disadvantages. The original definition may over prescribe, even for customers. In a preferred source, a lack of side effects may appear, and these may be necessary in programs. When you're developing to emulate the IBM 1401 on System 360, for example, I developed but there were 30 different "curios" side effects that spoiled my development, that had come into widespread use and had to be considered as part of the definition.

A *de facto* definition will often be found to be inadequate at these particularly precise locations, but it will be received as thought. This may cause will often turn out to be slow or slow to precisely in not be making it on. For example, some machines lack a shift register and a multiplier, and a multiplier. The precise nature of this trade is a part of the *de facto* definition, yet duplicating it may preclude the use of a faster multiplication algorithm.

In this chapter, we look at the work of several companies engaged in the game of DOS compatibility. These companies must figure out and use a reproduction of the deficient "curios" one might even say they had to use DOS definition. Since DOS, to my knowledge, is the DR-DOS, must of course be as close as possible to MS-DOS, we will provide additional features at a lower cost. In this chapter, we see how close Novell DOS comes to this goal.

Novell's NetWare is not a clone of DOS, but as we already saw briefly at Chapter 2, it hooks DOS and makes some important alterations to it. So discussing NetWare in more detail while we discuss



Novell DOS makes sense. In fact, we find that Novell NetWare makes so many changes to PC IN-286 interface that even though it seems to run "on top of" DOS, it effectively *replaces* DOS and thereby fully qualifies as an "other DOS."

IBM OS/2 2.2 and Windows NT most definitely do not hook or in any other way sit on top of DOS; they completely replace it. But market reality dictates that these environments run DOS programs out of the box. This ability is called *binary compatibility*; in contrast to the much simpler and older source compatibility. In fact, it seems that for some time to come, such environments planned for the purpose of running old DOS or Windows software. This requires that the new operating system either emulate DOS or run a copy of genuine DOS within a virtual machine. Either way, a important question is whether DOS programs that access undocumented DOS API structures (all undocumented DOS functions within these environments) find undocumented DOS a successful test of DOS compatibility. As having compatibility with undocumented interfaces stands as a so-called *you hear discussions of "DOS compatibility" or "Windows compatibility" or a variety of possible undocumented interfaces that's being discussed.*

An interesting point emerges from all this. It is that a *Microsoft interface* for many applications (not necessarily of your primary concern) to use undocumented DOS features as this ties how applications to Microsoft's own versions of DOS, more especially IBM's because of the DOS source code that is out of September 1993. After all, what's the real difference between Microsoft's DOS and anyone else's, except that Microsoft has guaranteed better support for the link and documented things that DOS applications do. For Microsoft, undocumented DOS within an interesting form of product differentiation.

## From CP/M to DR DOS to Novell DOS

The funny thing is MS-DOS itself started out as a clone of the CP/M operating system from DRI. The story has been told many times of how Intel did not want a computer manufacturer (at the first edition of this book, now at Microsoft since 1980) to give Intel Security Computer Products—two manufacturers Quick and Dirty DOS (QDOS)—now Intel says, No DOS, would a Microsoft purchase Intel's own rights and how this became MS-DOS for the then-new IBM PC. See also Microsoft and Paul Andrews' history of Microsoft. *Access via the catalog containing photos, pp. 47-50 or 211-214. See also computer order forms for Microsoft's purchase—No DOS, also Intel's Intel-550100.*

Quick, dirty, and cheap. As Andrew's *computer notes* means superbly text book on *Multi-Operating Systems*. It says clearly that what started out as a system that was never finished by accident was going to be controlling 50 million computers, considerably more than it might have gone into.

Somewhat understandably, Digital Research was upset when it found that Microsoft's new operating system for the IBM PC was a clone of CP/M. Apparently Digital's care could even considered string IBM over the semantics of MS-DOS. CP/M. Microsoft would be similarly upset by a of someone came out with a superior version that happened to provide the same APIs as Windows.

There's still a question, come MS-DOS's huge scale borrowing from CP/M. As Tom Patterson would write, somewhat later in *A Field Guide to MS-DOS*, *Byte* 1 in 1983. The primary design requirement of MS-DOS was to be MS-DOS translation compatible.

An early article in *Digital Contents*, "CP/M 86 vs. MS-DOS: A Technical Comparison," *Digital Journal* Jan. 1982, compared MS-DOS with both CP/M 80 on the Intel 8080 and CP/M 86 on the 8086, focusing not only where MS-DOS properly emulated CP/M functions but also where there were differences. For example, *Byte* 9, 1981, its original comparison with OS character 1 both systems, but CP/M expanded tabs and MS-DOS didn't. For example, by *System* is a reminder of MS-DOS's CP/M roots. To this day, MS-DOS contains many holdovers from its early start as a CP/M clone. The PSP, for example, is nothing more than a CP/M save page.

However, even at the beginning, there were crucial differences between the two systems. MS DOS did not allow a user to command a program to run every last C.P.M. function call. For example, MS DOS did not implement C.P.M. function 12 (R) to change the system's session number. Somewhat accidentally, MS DOS instead used a standard BIOS function 0C to read the keyboard.

There were other differences as well. In an important departure from the C.P.M. file system, MS DOS internally used a file allocation table (FAT) scheme borrowed from Microsoft's standard operating system. Paterson's original goal was to take the FAT memory resident at all times (in saving the time and disk reads that C.P.M. often made just to find where a file's data was located). As Ray Duvwan notes in the introduction to *16-Bit Software: The New Frontier* (November 1982), "Changes in operating systems appear very similar to the casual user; they are fundamentally different allocation schemes for the underlying files. This has significant, though subtle, effects on the speed of disk operations."

As MS DOS began to displace and enhance some of C.P.M.'s Digital Research makers of C.P.M., went to a "second generation" operating system, such as Concurrent C.P.M., EKSOS, GEM, or Microsoft ROM. Microsoft DOS and finally DR DOS. DR DOS was intended as a complete replacement for MS DOS. It was, however, never fully recognized by the original critics.

DR DOS now shows its C.P.M. roots. For one thing, the DR DOS disks and manuals carry copyright dates going back to 1976. Some of the code in DR DOS may even go back to the original C.P.M. code base. And some code was "re-used" C.P.M. terms such as PIP or CCP in the DR DOS manuals and code. DR DOS debugger, unfortunately, is not called DDD; the DR DOS kernel is still called BIOS, just as in olden days.

For example, the DR DOS SYSDISK driver has a BIOS option to relocate the DR DOS kernel to the GMA or to a 512 MB file-based disk option is available. So that MS DOS 5.0's DPMI-HELI continued to be a popular name. With providing more features than Microsoft's version of Microsoft DR DOS, it has a lot more C.P.M. heritage, at least in its naming conventions.

None of this would matter very much, except for the fact that Novell acquired Digital Research in 1980, and in July 1991, DR DOS now carries Novell's name. DR DOS may share in some of the success of Novell's NetWare. Novell's dominance in the network operating system market, far surpassing any of Microsoft's competitors, is due to providing network software. Novell's purchase of DR DOS shows why it was a "re-use" mistake, since DR DOS sales have dropped fairly sharply since the Novell purchase. However, development did slow down to Novell's purchase of DR DOS to Microsoft's release of MS DOS 5.0. DR DOS leads in price (as of the summer of 1990) at \$49.95, while Digital and DR DOS's standard Microsoft hardware is the venerable MS DOS version 4, engineered largely by IBM.

DR DOS's success in the marketplace for MS DOS, which is a lot about Microsoft's role since DR DOS was never a standard DR attempt to do a good job from 1,500,000 copies of DR DOS 5.0 in 1991, indicating that its success was also based on the great popularity of MS DOS 4.0. Novell's success in DR DOS comes from its 11% of the DOS market. *Winnipeg* (April 27, 1993) says the DR DOS's sales is 5% in 1991, while Microsoft sold 8.5 million worth of MS DOS. DR's total sales is over \$15 million. Assuming that DR's only marketing for DOS, and assuming that DR's sales are 10% of DR DOS, and MS DOS's sales are 10% of DR DOS, then the DOS market DR DOS's presence is only a small fraction of the United States, possibly in part because DR DOS was not of Novell's European Development Centre in Haverhill, England.

DR DOS does not give the impression of being very sporty, and more than one reviewer of DR DOS's code has been seen, however, to look at the code. The first piece of code, which is a DR DOS's code, is a MS DOS's code, and one reviewer, a former Microsoft employee, "showed the idea of any one would be checking on the presence of his other computer" and finally told us "I see so many of the same things that I see in the DR code, but they are programs that they developed a lot of their own code." This reflects a general feeling that DR DOS is based on and pretty much relevant for example, in an effort to optimize such choices for the 1990s. *The Magazine* (January 15, 1991), Charles Petzold placed DR DOS in the "Interesting But Does It Really Matter Department."

But consider the economics of DOS. For example, in its 1991 fiscal year (which ended in 1991), Microsoft had total revenues of \$1.843 billion. Of this, Windows (though not Windows applications like Excel and Word) brought in \$700 million (41%). Meanwhile, MS-DOS OEM price retail sales (again, no applications) brought—more than twice that total—\$617.5 million, which represents a whopping 33% of Microsoft's revenues.

Now notice that if a company takes only 5% of this business from Microsoft, it has made \$30 million. This could run a 300-person company on sales like those. As noted earlier, in 1991 DRI had total revenues of \$45 million. And the firm had 200 employees.

Why this discrepancy? Novell DOS is still alternative to MS-DOS and, coming with a support kit, this is not only because Novell DOS is a better implementation of DOS and because Novell DOS has stronger ties to NetWare, but also because Novell DOS is a somewhat more complete, more complete, with a lot of nice, nice aspects of MS-DOS that were either DR-DOS's or 6.0's.

Why would anyone buy DR-DOS rather than MS-DOS? One reason is the price of DR-DOS to computer manufacturers, who must handle a couple of DOS with each machine. Microsoft makes the OEM price \$10 and \$25 for the OEM copies of MS-DOS. DR-DOS's OEM price is lower—\$8 and \$15. And, as you can see, the OEM price of Novell DOS is even lower—\$5 for DR-DOS. Thus a manufacturer might consider handling DR-DOS rather than MS-DOS for the same reason that they might consider using a 386 chip from AMD or Cyrix rather than from Intel.

On the other hand, it is easy to see Microsoft's hegemony that it is next to impossible to find a machine loaded with DR-DOS rather than MS-DOS. For the simple reason that Microsoft has no competition in this market, which is effectively marketed for the copies of DOS loaded on every single PC. This is not a good thing for Microsoft's OEM pricing of MS-DOS since the machine rather than a per-copy basis. As Microsoft's Andrew notes in his statement of defense of MS-DOS OEM pricing, "per machine charges make it virtually impossible to crack the DOS monopoly. If we were simply paying a royalty to DOS on every machine made, your retailer likely to offer a better operating system, except for a high price. (quoting p. 26)

Another possibility is that DR-DOS might be that DR-DOS generally leads MS-DOS in features. The time line in Table 4-1 makes this clear.

**Table 4-1: Novell and DRI vs. Microsoft. A Timeline**

August 1990	DR-DOS 5.0, HI-DOS-OS, etc.
June 1991	MS-DOS 5.0, DOS-III, etc.
July 1991	Novell purchase of Digital Research
September 1991	DR-DOS 6.0 and compression kit
September 1991	NetWare 4.1
October 1992	Windows for Workgroups 3.1
April 1993	MS-DOS 6.0, DoubleSpace, etc.
April 1994	Acquisition of Novell DOS 7.0 and NetWare 4.1

DR-DOS and 6.0 had DOS key features like a better security, better file system, MS-DOS and it came bundled with disk compression, SuperDisk, and many other features. In any case, before MS-DOS, Microsoft's OEM price for Windows for Workgroups, WfW, some months after the release of Novell and DRI were bundled with NetWare, though both WfW and NetWare. The big fat head of Andrew's LAN work in the past to peer networking, DR-DOS appears to have been ROM-based, unlike MS-DOS. And what is perhaps the most important innovation of MS-DOS 5.0 was some variability to consumers, rather than only to OEMs for bundling with their machines—also concerning DR-DOS. DR-DOS 5.0 was the first retail DOS.

Basically, if you want to find out what features the next version of MS-DOS might support, you can look at the current version of DR-DOS. However, it is not necessarily true that Microsoft's 1994

copies dies from DR DOS. For that to be true, there would need to be more of a time lag between each DR DOS release and the subsequent release of the MS DOS version with similar features.

Most software has been released of carefully orchestrated, bulked-out later versions of MS DOS (i.e., an attempt to create what in the industry is called "E.D." fear uncertainty and doubt) regarding DR DOS. For example, in October, 1990 shortly after the release of DR DOS 5.0 and long before the execution of the 1991 version of MS DOS 5.0, stories on feature enhancements in MS DOS started to appear. *Forbes* (1) and *IT Week* (lead story) (2) Vice President of Systems Software at Microsoft and General Manager of Windows and MS DOS Business Unit, wrote a thoughtful letter to PC Week Newsletter (3), 1990, denoting that Microsoft was engaged in "E.D." tactics, "to serve our customers better, we will be more forthcoming about version 5.0" and denoting that Microsoft copies features from DR DOS. The feature enhancements of MS DOS version 5.0 were decided and developed before the only source we heard about DR DOS 5.0 (there will be some similar features). With standard MS DOS users, it shouldn't be surprising that DRJ has heard some of the same requests from customers that we have (4).

As a consumer, we don't get away, what is so bad about Microsoft's trying to patch features and to create a DR DOS. We should be thankful to DR DOS for whatever little competition it gives MS DOS and for whatever pressure it puts on Microsoft to keep improving DOS. The improvements of standards are making the DOS iterations through as it is. The present stagnation of DOS version 5.0 compared to other operating conditions of near monopolies. But with its planned competition in the system (DOS 7.0, Windows 4.0), Microsoft will, hopefully, make some of the expected improvements to DOS and caption was past Novell.

We see the similarities between standard DR DOS rather than MS DOS. In making such a comparison, important to remember that DR DOS is not MS DOS. As we see in this chapter, DR DOS has some differences not found in MS DOS but also has many compatibility differences in both known and undocumented areas of the DOS interface.

Before Novell DOS 5.0, a more complete of an operating system, Novell has reworked many of DR DOS' already made a new Novell DOS product is now much more compatible with MS DOS (improvement) (5) (6) (7) (8) (9) (10) (11) (12) (13) (14) (15) (16) (17) (18) (19) (20) (21) (22) (23) (24) (25) (26) (27) (28) (29) (30) (31) (32) (33) (34) (35) (36) (37) (38) (39) (40) (41) (42) (43) (44) (45) (46) (47) (48) (49) (50) (51) (52) (53) (54) (55) (56) (57) (58) (59) (60) (61) (62) (63) (64) (65) (66) (67) (68) (69) (70) (71) (72) (73) (74) (75) (76) (77) (78) (79) (80) (81) (82) (83) (84) (85) (86) (87) (88) (89) (90) (91) (92) (93) (94) (95) (96) (97) (98) (99) (100) (101) (102) (103) (104) (105) (106) (107) (108) (109) (110) (111) (112) (113) (114) (115) (116) (117) (118) (119) (120) (121) (122) (123) (124) (125) (126) (127) (128) (129) (130) (131) (132) (133) (134) (135) (136) (137) (138) (139) (140) (141) (142) (143) (144) (145) (146) (147) (148) (149) (150) (151) (152) (153) (154) (155) (156) (157) (158) (159) (160) (161) (162) (163) (164) (165) (166) (167) (168) (169) (170) (171) (172) (173) (174) (175) (176) (177) (178) (179) (180) (181) (182) (183) (184) (185) (186) (187) (188) (189) (190) (191) (192) (193) (194) (195) (196) (197) (198) (199) (200) (201) (202) (203) (204) (205) (206) (207) (208) (209) (210) (211) (212) (213) (214) (215) (216) (217) (218) (219) (220) (221) (222) (223) (224) (225) (226) (227) (228) (229) (230) (231) (232) (233) (234) (235) (236) (237) (238) (239) (240) (241) (242) (243) (244) (245) (246) (247) (248) (249) (250) (251) (252) (253) (254) (255) (256) (257) (258) (259) (260) (261) (262) (263) (264) (265) (266) (267) (268) (269) (270) (271) (272) (273) (274) (275) (276) (277) (278) (279) (280) (281) (282) (283) (284) (285) (286) (287) (288) (289) (290) (291) (292) (293) (294) (295) (296) (297) (298) (299) (300) (301) (302) (303) (304) (305) (306) (307) (308) (309) (310) (311) (312) (313) (314) (315) (316) (317) (318) (319) (320) (321) (322) (323) (324) (325) (326) (327) (328) (329) (330) (331) (332) (333) (334) (335) (336) (337) (338) (339) (340) (341) (342) (343) (344) (345) (346) (347) (348) (349) (350) (351) (352) (353) (354) (355) (356) (357) (358) (359) (360) (361) (362) (363) (364) (365) (366) (367) (368) (369) (370) (371) (372) (373) (374) (375) (376) (377) (378) (379) (380) (381) (382) (383) (384) (385) (386) (387) (388) (389) (390) (391) (392) (393) (394) (395) (396) (397) (398) (399) (400) (401) (402) (403) (404) (405) (406) (407) (408) (409) (410) (411) (412) (413) (414) (415) (416) (417) (418) (419) (420) (421) (422) (423) (424) (425) (426) (427) (428) (429) (430) (431) (432) (433) (434) (435) (436) (437) (438) (439) (440) (441) (442) (443) (444) (445) (446) (447) (448) (449) (450) (451) (452) (453) (454) (455) (456) (457) (458) (459) (460) (461) (462) (463) (464) (465) (466) (467) (468) (469) (470) (471) (472) (473) (474) (475) (476) (477) (478) (479) (480) (481) (482) (483) (484) (485) (486) (487) (488) (489) (490) (491) (492) (493) (494) (495) (496) (497) (498) (499) (500) (501) (502) (503) (504) (505) (506) (507) (508) (509) (510) (511) (512) (513) (514) (515) (516) (517) (518) (519) (520) (521) (522) (523) (524) (525) (526) (527) (528) (529) (530) (531) (532) (533) (534) (535) (536) (537) (538) (539) (540) (541) (542) (543) (544) (545) (546) (547) (548) (549) (550) (551) (552) (553) (554) (555) (556) (557) (558) (559) (560) (561) (562) (563) (564) (565) (566) (567) (568) (569) (570) (571) (572) (573) (574) (575) (576) (577) (578) (579) (580) (581) (582) (583) (584) (585) (586) (587) (588) (589) (590) (591) (592) (593) (594) (595) (596) (597) (598) (599) (600) (601) (602) (603) (604) (605) (606) (607) (608) (609) (610) (611) (612) (613) (614) (615) (616) (617) (618) (619) (620) (621) (622) (623) (624) (625) (626) (627) (628) (629) (630) (631) (632) (633) (634) (635) (636) (637) (638) (639) (640) (641) (642) (643) (644) (645) (646) (647) (648) (649) (650) (651) (652) (653) (654) (655) (656) (657) (658) (659) (660) (661) (662) (663) (664) (665) (666) (667) (668) (669) (670) (671) (672) (673) (674) (675) (676) (677) (678) (679) (680) (681) (682) (683) (684) (685) (686) (687) (688) (689) (690) (691) (692) (693) (694) (695) (696) (697) (698) (699) (700) (701) (702) (703) (704) (705) (706) (707) (708) (709) (710) (711) (712) (713) (714) (715) (716) (717) (718) (719) (720) (721) (722) (723) (724) (725) (726) (727) (728) (729) (730) (731) (732) (733) (734) (735) (736) (737) (738) (739) (740) (741) (742) (743) (744) (745) (746) (747) (748) (749) (750) (751) (752) (753) (754) (755) (756) (757) (758) (759) (760) (761) (762) (763) (764) (765) (766) (767) (768) (769) (770) (771) (772) (773) (774) (775) (776) (777) (778) (779) (780) (781) (782) (783) (784) (785) (786) (787) (788) (789) (790) (791) (792) (793) (794) (795) (796) (797) (798) (799) (800) (801) (802) (803) (804) (805) (806) (807) (808) (809) (810) (811) (812) (813) (814) (815) (816) (817) (818) (819) (820) (821) (822) (823) (824) (825) (826) (827) (828) (829) (830) (831) (832) (833) (834) (835) (836) (837) (838) (839) (840) (841) (842) (843) (844) (845) (846) (847) (848) (849) (850) (851) (852) (853) (854) (855) (856) (857) (858) (859) (860) (861) (862) (863) (864) (865) (866) (867) (868) (869) (870) (871) (872) (873) (874) (875) (876) (877) (878) (879) (880) (881) (882) (883) (884) (885) (886) (887) (888) (889) (890) (891) (892) (893) (894) (895) (896) (897) (898) (899) (900) (901) (902) (903) (904) (905) (906) (907) (908) (909) (910) (911) (912) (913) (914) (915) (916) (917) (918) (919) (920) (921) (922) (923) (924) (925) (926) (927) (928) (929) (930) (931) (932) (933) (934) (935) (936) (937) (938) (939) (940) (941) (942) (943) (944) (945) (946) (947) (948) (949) (950) (951) (952) (953) (954) (955) (956) (957) (958) (959) (960) (961) (962) (963) (964) (965) (966) (967) (968) (969) (970) (971) (972) (973) (974) (975) (976) (977) (978) (979) (980) (981) (982) (983) (984) (985) (986) (987) (988) (989) (990) (991) (992) (993) (994) (995) (996) (997) (998) (999) (1000) (1001) (1002) (1003) (1004) (1005) (1006) (1007) (1008) (1009) (1010) (1011) (1012) (1013) (1014) (1015) (1016) (1017) (1018) (1019) (1020) (1021) (1022) (1023) (1024) (1025) (1026) (1027) (1028) (1029) (1030) (1031) (1032) (1033) (1034) (1035) (1036) (1037) (1038) (1039) (1040) (1041) (1042) (1043) (1044) (1045) (1046) (1047) (1048) (1049) (1050) (1051) (1052) (1053) (1054) (1055) (1056) (1057) (1058) (1059) (1060) (1061) (1062) (1063) (1064) (1065) (1066) (1067) (1068) (1069) (1070) (1071) (1072) (1073) (1074) (1075) (1076) (1077) (1078) (1079) (1080) (1081) (1082) (1083) (1084) (1085) (1086) (1087) (1088) (1089) (1090) (1091) (1092) (1093) (1094) (1095) (1096) (1097) (1098) (1099) (1100) (1101) (1102) (1103) (1104) (1105) (1106) (1107) (1108) (1109) (1110) (1111) (1112) (1113) (1114) (1115) (1116) (1117) (1118) (1119) (1120) (1121) (1122) (1123) (1124) (1125) (1126) (1127) (1128) (1129) (1130) (1131) (1132) (1133) (1134) (1135) (1136) (1137) (1138) (1139) (1140) (1141) (1142) (1143) (1144) (1145) (1146) (1147) (1148) (1149) (1150) (1151) (1152) (1153) (1154) (1155) (1156) (1157) (1158) (1159) (1160) (1161) (1162) (1163) (1164) (1165) (1166) (1167) (1168) (1169) (1170) (1171) (1172) (1173) (1174) (1175) (1176) (1177) (1178) (1179) (1180) (1181) (1182) (1183) (1184) (1185) (1186) (1187) (1188) (1189) (1190) (1191) (1192) (1193) (1194) (1195) (1196) (1197) (1198) (1199) (1200) (1201) (1202) (1203) (1204) (1205) (1206) (1207) (1208) (1209) (1210) (1211) (1212) (1213) (1214) (1215) (1216) (1217) (1218) (1219) (1220) (1221) (1222) (1223) (1224) (1225) (1226) (1227) (1228) (1229) (1230) (1231) (1232) (1233) (1234) (1235) (1236) (1237) (1238) (1239) (1240) (1241) (1242) (1243) (1244) (1245) (1246) (1247) (1248) (1249) (1250) (1251) (1252) (1253) (1254) (1255) (1256) (1257) (1258) (1259) (1260) (1261) (1262) (1263) (1264) (1265) (1266) (1267) (1268) (1269) (1270) (1271) (1272) (1273) (1274) (1275) (1276) (1277) (1278) (1279) (1280) (1281) (1282) (1283) (1284) (1285) (1286) (1287) (1288) (1289) (1290) (1291) (1292) (1293) (1294) (1295) (1296) (1297) (1298) (1299) (1300) (1301) (1302) (1303) (1304) (1305) (1306) (1307) (1308) (1309) (1310) (1311) (1312) (1313) (1314) (1315) (1316) (1317) (1318) (1319) (1320) (1321) (1322) (1323) (1324) (1325) (1326) (1327) (1328) (1329) (1330) (1331) (1332) (1333) (1334) (1335) (1336) (1337) (1338) (1339) (1340) (1341) (1342) (1343) (1344) (1345) (1346) (1347) (1348) (1349) (1350) (1351) (1352) (1353) (1354) (1355) (1356) (1357) (1358) (1359) (1360) (1361) (1362) (1363) (1364) (1365) (1366) (1367) (1368) (1369) (1370) (1371) (1372) (1373) (1374) (1375) (1376) (1377) (1378) (1379) (1380) (1381) (1382) (1383) (1384) (1385) (1386) (1387) (1388) (1389) (1390) (1391) (1392) (1393) (1394) (1395) (1396) (1397) (1398) (1399) (1400) (1401) (1402) (1403) (1404) (1405) (1406) (1407) (1408) (1409) (1410) (1411) (1412) (1413) (1414) (1415) (1416) (1417) (1418) (1419) (1420) (1421) (1422) (1423) (1424) (1425) (1426) (1427) (1428) (1429) (1430) (1431) (1432) (1433) (1434) (1435) (1436) (1437) (1438) (1439) (1440) (1441) (1442) (1443) (1444) (1445) (1446) (1447) (1448) (1449) (1450) (1451) (1452) (1453) (1454) (1455) (1456) (1457) (1458) (1459) (1460) (1461) (1462) (1463) (1464) (1465) (1466) (1467) (1468) (1469) (1470) (1471) (1472) (1473) (1474) (1475) (1476) (1477) (1478) (1479) (1480) (1481) (1482) (1483) (1484) (1485) (1486) (1487) (1488) (1489) (1490) (1491) (1492) (1493) (1494) (1495) (1496) (1497) (1498) (1499) (1500) (1501) (1502) (1503) (1504) (1505) (1506) (1507) (1508) (1509) (1510) (1511) (1512) (1513) (1514) (1515) (1516) (1517) (1518) (1519) (1520) (1521) (1522) (1523) (1524) (1525) (1526) (1527) (1528) (1529) (1530) (1531) (1532) (1533) (1534) (1535) (1536) (1537) (1538) (1539) (1540) (1541) (1542) (1543) (1544) (1545) (1546) (1547) (1548) (1549) (1550) (1551) (1552) (1553) (1554) (1555) (1556) (1557) (1558) (1559) (1560) (1561) (1562) (1563) (1564) (1565) (1566) (1567) (1568) (1569) (1570) (1571) (1572) (1573) (1574) (1575) (1576) (1577) (1578) (1579) (1580) (1581) (1582) (1583) (1584) (1585) (1586) (1587) (1588) (1589) (1590) (1591) (1592) (1593) (1594) (1595) (1596) (1597) (1598) (1599) (1600) (1601) (1602) (1603) (1604) (1605) (1606) (1607) (1608) (1609) (1610) (1611) (1612) (1613) (1614) (1615) (1616) (1617) (1618) (1619) (1620) (1621) (1622) (1623) (1624) (1625) (1626) (1627) (1628) (1629) (1630) (1631) (1632) (1633) (1634) (1635) (1636) (1637) (1638) (1639) (1640) (1641) (1642) (1643) (1644) (1645) (1646) (1647) (1648) (1649) (1650) (1651) (1652) (1653) (1654) (1655) (1656) (1657) (1658) (1659) (1660) (1661) (1662) (1663) (1664) (1665) (1666) (1667) (1668) (1669) (1670) (1671) (1672) (1673) (1674) (1675) (1676) (1677) (1678) (1679) (1680) (1681) (1682) (1683) (1684) (1685) (1686) (1687) (1688) (1689) (1690) (1691) (1692) (1693) (1694) (1695) (1696) (1697) (1698) (1699) (1700) (1701) (1702) (1703) (1704) (1705) (1706) (1707) (1708) (1709) (1710) (1711) (1712) (1713) (1714) (1715) (1716) (1717) (1718) (1719) (1720) (1721) (1722) (1723) (1724) (1725) (1726) (1727) (1728) (1729) (1730) (1731) (1732) (1733) (1734) (1735) (1736) (1737) (1738) (1739) (1740) (1741) (1742) (1743) (1744) (1745) (1746) (1747) (1748) (1749) (1750) (1751) (1752) (1753) (1754) (1755) (1756) (1757) (1758) (1759) (1760) (1761) (1762) (1763) (1764) (1765) (1766) (1767) (1768) (1769) (1770) (1771) (1772) (1773) (1774) (1775) (1776) (1777) (1778) (1779) (1780) (1781) (1782) (1783) (1784) (1785) (1786) (1787) (1788) (1789) (1790) (1791) (1792) (1793) (1794) (1795) (1796) (1797) (1798) (1799) (1800) (1801) (1802) (1803) (1804) (1805) (1806) (1807) (1808) (1809) (1810) (1811) (1812) (1813) (1814) (1815) (1816) (1817) (1818) (1819) (1820) (1821) (1822) (1823) (1824) (1825) (1826) (1827) (1828) (1829) (1830) (1831) (1832) (1833) (1834) (1835) (1836) (1837) (1838) (1839) (1840) (1841) (1842) (1843) (1844) (1845) (1846) (1847) (1848) (1849) (1850) (1851) (1852) (1853) (1854) (1855) (1856) (1857) (1858) (1859) (1860) (1861) (1862) (1863) (1864) (1865) (1866) (1867) (1868) (1869) (1870) (1871) (1872) (1873) (1874) (1875) (1876) (1877) (1878) (1879) (1880) (1881) (1882) (1883) (1884) (1885) (1886) (1887) (1888) (1889) (1890) (1891) (1892) (1893) (1894) (1895) (1896) (1897) (1898) (1899) (1900) (1901) (1902) (1903) (1904) (1905) (1906) (1907) (1908) (1909) (1910) (1911) (1912) (1913) (1914) (1915) (1916) (1917) (1918) (1919) (1920) (1921) (1922) (1923) (1924) (1925) (1926) (1927) (1928) (1929) (1930) (1931) (1932) (1933) (1934) (1935) (1936) (1937) (1938) (1939) (1940) (1941) (1942) (1943) (1944) (1945) (1946) (1947) (1948) (1949) (1950) (1951) (1952) (1953) (1954) (1955) (1956) (1957) (1958) (1959) (1960) (1961) (1962) (1963) (1964) (1965) (1966) (1967) (1968) (1969) (1970) (1971) (1972) (1973) (1974) (1975) (1976) (1977) (1978) (1979) (1980) (1981) (1982) (1983) (1984) (1985) (1986) (1987) (1988) (1989) (1990) (1991) (1992) (1993) (1994) (1995) (1996) (1997) (1998) (1999) (2000) (2001) (2002) (2003) (2004) (2005) (2006) (2007) (2008) (2009) (2010) (2011) (2012) (2013) (2014) (2015) (2016) (2017) (2018) (2019) (2020) (2021) (2022) (2023) (2024) (2025) (2026) (2027) (2028) (2029) (2030) (2031) (2032) (2033) (2034) (2035) (2036) (2037) (2038) (2039) (2040) (2041) (2042) (2043) (2044) (2045) (2046) (2047) (2048) (2049) (2050) (2051) (2052) (2053) (2054) (2055) (2056) (2057) (2058) (2059) (2060) (2061) (2062) (2063) (2064) (2065) (2066) (2067) (2068) (2069) (2070) (2071) (2072) (2073) (2074) (2075) (2076) (2077) (2078) (2079) (2080) (2081) (2082) (2083) (2084) (2085) (2086) (2087) (2088) (2089) (2090) (2091) (2092) (2093) (2094) (2095) (2096) (2097) (2098) (2099) (2100) (2101) (2102) (2103) (2104) (2105) (2106) (2107) (2108) (2109) (2110) (2111) (2112) (2113) (2114) (2115) (2116) (2117) (2118) (2119) (21



**Table 4-2: DR DOS INT 21h functions**

AX=4302h	Get Access Rights (documented in DR DOS 6.0)
AX=4305h	Set Access Rights and Password (documented)
AX=4304h	Get Encrypted Password
AX=4308h	Set Extended File Attributes
AX=4306h	Get File Owner
AX=4307h	Set File Owner
AX=4380h	Lock/Unlock Pending Delete file (Novell DOS 7.0)
AX=4381h	Purge Pending Delete file (Novell DOS 7.0)
AX=4451h	Concurrent DOS Install Check
AX=4452h	Get DR DOS Version
AX=4454h	Set Global Password
AX=4456h	DR DOS History Buffer Control
AX=4457h	Get/Set Share Thread Status
AX=4458h	Get Pointer to Internal Variable Table
AX=4100h	(CX=88h) Find First Deleted File
AX=4101h	(CX=88h) Find Next Deleted File

For example, the value returned by INT 21h AX=4452h (Get DR DOS Version) is probably not correct for DOS 6.04. In 20 is the ASCII code for the number 10. Under DR DOS and Novell DOS, this is the same as the DR product code in AX and in Novell DOS 7.0, release code in DX. Elsewhere, such as under MS-DOS, this instruction should return with the carry flag set and/or with the value in AX=4452h in DR DOS-based products such as Concurrent DOS and Multuser DOS use a similar function, INT 21h AX=4451h.

The instruction INT 21h DR DOS (not on method) has unfortunately been a closely guarded DR DOS secret. Some computers reportedly found the DR DOS (the official DOS) in Germany. "I just had a long DR DOS phone call. The phone was asked for files, it was a waste, I don't have DR DOS. I had to get my files, I don't have the files, I don't have MS-DOS 3.31, not DR DOS 5.0, are absolutely the same."

Novell's security-conscious business policy of pretending that DR DOS is identical to MS-DOS is a good example of official corporate support for DR DOS. *System and Programmer's Guide* (DR product #1182 2013 001) the very existence of which is sometimes even denied. DR DOS programmers' common knowledge is that DOS 6.04 is 100% ready available through Novell's developer relations centers in Austin, TX. The only case where documentation a few DR DOS-specific functions such as AX=4307h and AX=4305h (see Table 4-2) and the TaskMAX INT 21h AH=27h interface, they may not even mention the single piece of information that most developers want about DR DOS, which is how to detect that you're running under it.

The DR DOS version codes (not simply DOS version numbers) instead they indicate the code for a DR DOS product release numbers and the way back to function 12 in CP/M 80. Table 4-3 shows the key version numbers.

**Table 4-3: 1067h and All That: Novell/DR DOS Product Codes**

1067h	DR DOS 3.41
1065h	DR DOS 5.0
1067h	DR DOS 6.0
1070h	DR PalmDOS
1071h	DR DOS 6.4, March 1993, publication W/W
1072h	Novell DOS 7.0
1466h	DR Multuser DOS 5.1
1467h	Concurrent DOS 5.1

As you can see, single user (through possibly multitasking, such as Novell DOS 3.0) operating systems use `10h` and multiuser systems use `14h`. Function `4452h` only works under the single user systems such as DR-DOS, and `4451h` works only under the multiuser systems such as Concurrent DOS, DR 3.11, EDOS, or OEM versions of DR-DOS on old desktop computers. Many of the DOS compatibility improvements found in Novell DOS 7.0 were in fact first introduced in early DOS. The DOS compatibility kernel is also present in the March 1993 "It's new, update" of DR-DOS so programs that need the new BIOS can check for `1070h` or higher (that is, `AI > 1070` or `AI == 10h` and `AI == 70h`).

All this can be wrapped in a C callable function to custom show (see Fig 4-1) in `DRDOS.C`. You can incorporate these checks into programs or, by compiling it with `LINK /NDALONE`, run it as a stand-alone test. You might contrast the straightforward nature of this code with the strangeness of `MSD-EJECT.C` in Chapter 1.

#### Listing 4-1: IS\_DRDOS.C

```

/*
IS_DRDOS.C
Andrew Schulman, January 1993
With changes by John Frotz, Realism of Novell, February 1993
From "Undocumented DOS", 2nd edition (Addison-Wesley, 1993)
To link with other programs: bcc -c is_drdos.c
To make standalone test: bcc -OSTANDALONE %s drdos.c
*/

#include <stdib.h>
#include <stdio.h>
#include <dos.h>

typedef unsigned short WORD;
/* EDC = European Development Centre, Novell Digital Research
Systems Group, Hungerford, England */
WORD EDC_product_codes WORD api_function;
{
    _asm mov ax, api_function
    _asm stc /* set carry flag */
    _asm int 21h
    _asm jc no_drDOS /* carry set: function not supported */
    _asm cmp ax, api_function
    _asm je no_drDOS /* AX unchanged, function not supported */
have_drDOS:
    _asm jmp short done /* carry clear and AX changed */
no_drDOS:
    _asm xor ax, ax /* not supported, return AX=0 */
done:
    /* return value in AX */
}

WORD is_drDOS(void)
{
    WORD product_code;

    /* Try the single-user API call */
    if ((product_code = EDC_product_codes(0x4452)) != 0)
        return product_code;

    /* Try the multi-user API call */
    if ((product_code = EDC_product_codes(0x4451)) != 0)
        return product_code;
}

/* If 1
/* The following seems like a bad idea, because the documented
call already go up to AX=447h (Query IOCTL Device) */
#else

```

```

/* Try an _old_and_obsolete_ single-user API call */
if (!product_code = EDC_product_codes(0x4412)) {
    return product_code;
}

/* still here - definitely no DR DOS! */
return 0;
}

#ifdef STANDALONE
main()
{
    WORD EDC_product;

    if ((EDC_product = %s_drdoes()) == 0)
        printf("Single- or Multi-user DR DOS not running\n");
    else
    {
        switch (EDC_product)
        {
            case 0x1060: printf("DOS Plus 1.1, break;"); break;
            case 0x1063: printf("DR DOS 5.41 1, break;"); break;
            case 0x1064: printf("DR DOS 5.42 1, break;"); break;
            case 0x1065: printf("DR DOS 5.0, break;"); break;
            case 0x1067: printf("DR DOS 6.0, break;"); break;
            case 0x1070: printf("DR DOS PatmDOS 1, break;"); break;
            case 0x1071: printf("DR DOS 6.0 March 1993 update"); break;
            case 0x1072: printf("Novell DOS 7.0"); break;
            case 0x1452: printf("Concurrent PC-DOS 3.2"); break;
            case 0x1461: printf("Concurrent DOS 4.1"); break;
            case 0x1450: printf("Concurrent DOS/386 5.0 or "
                "Concurrent DOS/386 1.1"); break;
            case 0x1460: printf("Concurrent DOS/386 6.0 or "
                "Concurrent DOS/386 2.0"); break;
            case 0x1462: printf("Concurrent DOS/386 6.2 or "
                "Concurrent DOS/386 3.0"); break;
            case 0x1466: printf("DR Multiuser DOS 5 1"); break;
            case 0x1467: printf("Concurrent DOS 5 1"); break;
            default: printf("unknown DR DOS version"); break;
        }

        printf("(B005 104KH)\n", EDC_product);
        printf("Providing DOS interface via IO2u\n", osmajor, osminor);
    }

    return (EDC_product);
}
#endif

```

IS DR/DOS has been tested under MS-DOS 5.0 and 6.0 where it correctly outputs "Single- or Multi-user DR DOS not running" and under DR-DOS 5.0 6.0, Novell DOS 7.0 and under Concurrent DOS 386 1.1 which are all systems purchased at the World Mail Warehouse in San Jose, CA home of dead hardware and software for \$14.95.

Of course, the cost of the user at home with DR-DOS and Novell DOS call function 4452h to write to the registry is not a DR-DOS. For example, if you take COMMAND.COM or MEM.MEM under DR-DOS 6.0 and run it on the machine under MS-DOS 6.0 they will complain "Incorrect system operation" etc. I don't think this is not done correctly. For example, the MEM.MEM file under DR-DOS 6.0 runs under MS-DOS and while it generally produces correct results, the /S shows where structure is questionable to ensure that it is indeed running under DR-DOS, running DR's MEM/S under MS-DOS hangs the system.



## Watching DR DOS

Assuming one already knows about the undocumented DR DOS functions, how does one go about determining which DR DOS utilities use them? The INTRSPY utility, presented in Chapter 5, is perfect for this. The DRDOS SCR script shown in Listing 4-2 lets INTRSPY watch some selected DR DOS functions. In addition to the INT 21h categories from Table 4-2, DRDOS SCR also reports on INT 2fh AA-12Fh, which the MEMMAX utility uses.

### Listing 4-2: INTRSPY Script for Selected DR DOS Functions

```

, DRDOS.SCR
intercept 2fh
  function 43h on_entry
    if (al > 1)
      output "21/43/" al
  function 44h on_entry
    if (al > 11h)
      output "21/44/" al
  function 4bh on_entry
    output (DS:DX->byte,ascii,64)
intercept 2fh
  function 12h
    subfunction 0FFh on_entry
      output "2F/12FF/" bx " /" cx " /" dx

```

Figure 4-1 shows a sample INTRSPY script while running the simple MEM and MEMMAX utilities in DR DOS 6.0.

### Figure 4-1: Undocumented DR DOS Calls from MEM and MAXMEM (DR DOS 6.0)

```

A \MEM.EXE
21/44/52
21/44/52
21/44/58
2F/12FF/0006/0000/0000
21/44/58
2F/12FF/0006/0000/0000
21/44/58
21/44/58
21/44/57
21/44/56
21/44/56
21/44/57
21/44/56
21/44/56
21/44/56
A \MEMMAX.EXE
2F/12FF/0006/0000/0000
2F/12FF/0006/0000/0000
2F/12FF/0006/0000/0000
21/44/57
21/44/56
21/44/56

```

To effectively use INTRSPY, however, you must already know what you are looking for. INTRSPY is good for checking if a MEM and MAXMEM call INT 21h AA-4452h, for example, but to write DRDOS SCR, you had to know already that DR DOS makes IOCHH calls with hexion numbers greater than 11h. How does one figure that out?

## Disassembling DR DOS

Through disassembly of the DR DOS files, of course. The system files in DR DOS include IBMBIO.COM and IBMDOS.COM. The command interpreter is, of course, COMMAND.COM.

Dbase III with a tool such as Source shows, for example, that COMMAND.COM calls INT 21h AX 4421h BX 4156h AX 4479h and AX 4458h. IBM DOS.COM provides these calls.

For some other disassemblers, the separate utilities shipped with DR DOS, such as MEMEAT and MEMMAX.EXE. However, Novell compresses these programs with PKLITE, an executable file compressor, into PKWARE. To disassemble the files, you must first decompress them with the PKLITE. A special PKLITE is readily available electronically, for example from the IBM SYS forum on CompuServe.

As noted earlier, programs such as MEMMAX.EXT, which in DR DOS enables, disables, and displays features of upper, lower, and video memory, together expands on what PKLITE. A MEMMAX.EXT version 7K (this version is what Source produces a small 100-line page listing in which it is easy to find the code) also studies the MEMMAX system, such as 1. to open upper memory for HIGHER\_U, 2. to use 1. A certain minimum free key memory space, 3. to use over the first 64K of memory, similar to HIGHER\_U in MS DOS, and so on.

Since MEMMAX.EXT is available, for example, how DR DOS enables and disables upper memory, the DR DOS 5.0 6.0 do not require such as MS DOS 5.0 M1-58h functions for upper memory, since Novell DOS 5.0 does. MEMMAX contains inline code to manipulate the Memory List of INT 21h M1-52, 53, 54, 55, pointer to SYSVARS, of which DR DOS, as we will see, makes upper memory and uses SYSVARS 7 to get a far pointer to the MCB chain.

Since INT 21h AX 4156h BX 4156h, MS DOS 5.0 and 6.0, MEMMAX enables upper memory, and so on, with the use of the MCBs, some instructions replacing the last MCB's Z flag with an M'. Later, in the next section, we describe upper memory MCB chain. To disable UMBs, it finds a last block with ADDR 0000, memory size M1 and a Z flag, simply come, see string 4-3 later in this chapter. Also, it finds a last block with a last MCB's address of INT 21h M1-58h function.

The parameters A, A, and A, and other memory switches MEMMAX does not use inline code instead of the instructions of the program. INT 21h AX 1211h BX 615h 0 with substantial information to CA.

```
0 get status of video memory space (memmax /v)
1 set memory into video memory space (memmax /u)
2 strip memory from video memory space (memmax /w)
```

use IN 21h AX 4156h BX 4156h inside HIGHER\_U

## How Close Is DR DOS to MS DOS?

DR DOS's price, low to MS DOS. MS DOS's services, such as low-level systems software, such as Windows and other utilities, are, of course, Novell's solution with DR DOS. The DR DOS 6.01, 6.02, 6.03, and 6.04, that is, 6.01, 6.02, 6.03, and 6.04, are generally compatible with applications that expect to run together, but they do have the standard conventions for DOS program standard conventions. These conventions include: support for larger than 32 MB partitions using the CDM (CD-ROM) file system, and 26-bit addresses. However,

Some applications would be designed to go beyond these conventions and actually attempt to change or alter the DOS file structures or replace sections of the operating system code with their own. This problem is independent on having intimate knowledge of each DOS version, but it is a fact that the software to react differently for each of these versions is a complex task. For this reason, the manufacturer will have to design the program to take the DR DOS 6.0 operating system into account as well.

For most users, it is a 50-50% of the application's responsibility to conform to DR DOS, but DR DOS's responsibility is to conform to the application because, as we noted earlier, that is how *de facto* standards work. In DR DOS, the manufacturer conforms to the conventions of DOS programming, not the other way around. Hence, the new Novell DOS 6.0 kernel.

Indeed, with each release DR DOS has incorporated support for more and more of the *de facto* DOS programming standards. DR DOS started out as a concurrent DOS. DR stripped out the concurrent and multi-tasking parts. It looked as if MS DOS compatibility. After adding to MS DOS compatibility was added with each release, as it seemed to be needed. One person feeling that DR had not supporting any of another undocumented DOS internal data structures, such as the Current Directory Structure, for as long as it could. This produced some compatibility problems, since as those are not earlier with Atari. In Novell DOS 7.0. However, or actually any DOS 11.70, there is finally a very close work alike to MS DOS 5.31. Note, however, that even these newer versions do not pass the highly arbitrary test made by Microsoft Windows VARD detection code, discussed in Chapter 2.

A technical support engineer at Novell reports that "many of our escalated technical support problems with developers are a result of differences that DR DOS 6.0 has with respect to the test queries used within *Unsubstantial DOS*." Presumably some of these problems are documented in the first edition of this book, but in some cases DR DOS 6.0 was simply not sufficient, especially with undocumented but widely used aspects of MS DOS.

Clearly, it would be unrealistic for Novell and other builders of DOS work-alikes if DOS applications would restrict themselves to the documented DOS interface. Seen in this light, the widespread use of undocumented DOS is really to Microsoft's advantage. Novell would eventually prefer to apply it only where it is needed. Undocumented DOS calls are a last resort. But the same technical support engineer just quoted also reports that at Novell, "most of our technical analysis that SeeWare File Fix was developed almost exclusively from *Unsubstantial DOS*." This clearly is a common situation. Many developers of non-PC software, especially because of the widespread use of undocumented DOS calls, must suffer, frankly, because of the widespread use of this book. Yet, at the same time, these companies are these early themselves. They just wish everyone would follow some restraint! Of course, Microsoft sets the standard here. With this, I say not as a deliberate attempt to undocumented calls.

### **SysVars, the Current Directory Structure, and the Redirector**

We've already noted that Novell's own MEMMAN and INITDOS programs rely on the core INT 21h and AH=52h undocumented functions to this extent, and at least some of its associated SysVars structure clearly may be present in DR DOS. Most of the programs in other parts of this book will run under DR DOS. Windows also runs on top of DR DOS, though, as noted several times already, for this to happen DR DOS 5 and to some extent to MS DOS 5.31. Clearly, DR DOS is a close approximation to even the undocumented portions of MS DOS.

However, under DR DOS 5.0 and 6.0 some of the critical data structures aside from undocumented DOS calls are *undocumented*. These structures exist to keep applications happy, and DR DOS does not actually use these structures. Applications with *quirks* rather than *bugs* expect these structures may be in for a rude shock.

For example, the `LASTDRIVE` value from Chapter 2 (modified by CDS pointer) and `LASTDRIVE` fields in `SysVars`. `LASTDRIVE` appears at first to contain the letter for DR DOS. Function 57 returns a pointer to `SysVars`, and `SysVars` contains what appears to be a valid CDS pointer and a `LASTDRIVE` value. `LASTDRIVE` successfully modifies both names. However, this has no effect on the system. If the value of `LASTDRIVE` in `CONFIG.SYS` was `F`, and if `LASTDRIVE` changes the `LASTDRIVE` value in `SysVars` to `F`, the highest valid drive is still `F`. Quarterback's `LASTDRIVE.COM` or `ls` from QMM However, has no effect on the number of actual drives under DR DOS 5.0 and 6.0.

The problem is that DR DOS 5.0 and 6.0 keep their genuine current directory and file information in software, not a suppliance enough of a CDS to satisfy programs that a CDS does indeed exist. DR DOS 5.0 and 6.0 internally store the current directory based on its first cluster number than on its ANSI name. DR DOS 6.0 rebuilds the fictitious CDS at appropriate times from DRI's



**Listing 4-3: Generic UMB Link Functions**

```

( from Arne Schapers, "DOS 5 for Programierer", pp 684-688 )
( Returns TRUE if UMBs are linked in, works for MS-DOS 5, MS-DOS 6.0,
DR-DOS 5.0, and DR-DOS 6.0 )
Function GetUMBLink: Boolean;
var Regs: Registers,
    MCB: MCBPtr; p: ^Word;
begin
  if TrueDosVersion = $3203 then ( DR-DOS 5.0 )
  begin
    ( GetDosDataArea calls INT 21h AH=52h )
    p := GetDosDataArea, Dec(LongInt(p),2);
    MCB := Ptr(p^,0), ( DOS Data Area, offset - 2 )
    while MCB^.Flag <> 'Z' do ( walk MCB chain )
      MCB := Ptr(Seg(MCB^)+1+MCB^.Size,0);
    GetUMBLink := Seg(MCB^)+MCB^.Size > $A000,
  end else
  begin
    Regs.AX := $5802, ( Function, "Get UMB Link State" )
    Intr($21,Regs);
    if Regs.Flags and FCarry <> 0 then GetUMBLink := False
    else GetUMBLink := (Regs.AL <> 0), ( TRUE if UMBs linked in )
  end,
end,
( Tries to link UMBs in, works with MS and DR-DOS 5.0/6.0 )
Procedure SetUMBLink( LinkState: Boolean);
var Regs: Registers,
    MCB: MCBPtr; p: ^Word,
    NextMCB: Word, Done: Boolean;
begin
  if DosVersion = $3203 then ( DR-DOS 5.0 )
  begin
    p := GetDosDataArea; Dec(LongInt(p),2);
    MCB := Ptr(p^,0), ( DOS Data Area, offset - 2 )
    Done := False,
    ( walk MCB chain to <= $9FFF )
    while (MCB^.Flag <> 'Z') and not Done do
    begin
      NextMCB := Seg(MCB^)+1+MCB^.Size,
      if NextMCB >= $9FFF then Done := True
      else MCB := Ptr(NextMCB,0);
    end;
    if NextMCB < $9fff then NextMCB := Seg(MCB^)+1+MCB^.Size,
    if (NextMCB < $9FFF) and ( MCB with ID of MEMMAX on $9FFF )
    ( MCBPtr(Ptr(NextMCB,0))^OwnerPSP = $0007) then
    begin
      if LinkState then MCB^.Flag := 'M'
      else MCB^.Flag := 'Z',
    end
  end else DosError := 7; ( "Function not supported" )
end else
begin
  ( MS-DOS )
  Regs.AX := $5803, ( Function "Set UMB Link" )
  Regs.BX := Ord(LinkState),
  Intr($21,Regs);
  if Regs.Flags and FCarry <> 0 then DosError := Regs.AX
  else DosError := 0,
end,
end;
end;

```

On DR DOS 5.0 and 6.0 require these replacements as Novell DOS 7.0 implements the DOS 5.0 INT 21h AX-5802h and AX-5803h functions.

### **TSRs and the Swappable Data Area**

DR DOS implements all the formerly undocumented functions commonly used by memory resident DOS programs. These functions include INT 21h AH=34h Get InDOS flag, AH=50h Set PSP, AH=51h Get PSP, AX=5100Ah Set Extended Error Information, and INT 28h Keyboard Idle. DOS TSRs run under DR DOS without incident.

DR DOS also implements undocumented DOS function INT 71h AX=5400h which returns a pointer to the Swappable Data Area. Some TSRs use this to swap the DOS state (see Chapter 9). Writing the handle in the SDI vector of the same way as MS DOS, most programs don't look at individual fields in the SDI. They just swap the entire structure using the swap EDI/DI and swap ESI/EAX sizes returned from function 5400h. This appears to work under all versions of DR DOS, though naturally the entire DOS state is not intact kept in the SDI. But this is true in MS DOS as well. Writing SDI swapping TSRs is risky business.

### **Additional DR DOS and Novel DOS Functionality**

Having looked at how DR DOS falls just slightly short of MS DOS in implementing various undocumented DOS functions and structures, let's take a moment to look at what *additional* functionality DR DOS provides for programmers.

The most important extra DR DOS functionality is of course the DR DOS detection call INT 21h AX=4457h. I discuss this elsewhere, but there are a number of other undocumented DR DOS functions, it seems unlikely that programmers would know about the DR DOS specific software (the only thing that makes sense to call this a DR DOS "troubleshooter" is a large market). It is more likely that programmers would want to take advantage of some undocumented interfaces provided in DR DOS and Novel DOS such as the DrWait, DrKick APIs, which Novell documents in a specification on "File and Directory Salvage for DOS Media."

The MAX file multitasking interface added to DR DOS 6.0 has a set of INT 21h AH=27h functions. I've covered this in the *DOS 6.05 Technical Programmer's Guide* (August 1991). INTRINS on the accompanying disk describes these calls. For example, INT 21h AX=2704h registers a new task on the Windows disk. It describes these calls. For example, INT 21h AX=2704h registers a new task on the Windows disk. For further information, see the DrWait, DrKick APIs, which Novell documents in a specification on "File and Directory Salvage for DOS Media."

In Novell DOS 6.0, KRNI 586 SYS provides true preemptive multitasking capabilities. If KRNI 386 SYS is installed, Novell DOS 6.0 also provides all MS services supporting either the 0-96 or 1-11 extension of the DPMI specification. All the DPMI API programs from Chapter 3 run under Novell DOS 6.0. If KRNI 586 SYS is installed, DPMI is supported.

Novell DOS 6.0 also introduces a new interface called the DOS Protected Mode Services (DPMS). A DPMS user's call documentation is available from Novell, as well as the Interrupt List on disk INT 21h AX=4310h.

The name DPMS is a mix of the DPMI and DOS file, the common American acronym for protection and force, but—as with the, to do with either, DPMS is a set of services that allow DOS 6.0s and device drivers to come themselves into extended memory and run in protected mode. The key behind DPMS is to make it very easy to port TSRs and device drivers to protected mode. For example, DPMS took a long time to write, but protected mode VxDs.

DPMS is significant because Novell can be using it for its own utilities, not only in Novell DOS 7.0, but also in NetWare Lite. Novell plans to use DPMS for disk cache and compression software, CD-ROM extensions, resident workstation management utilities, workstation shells, redirectors, and requests.

## Protected Mode DOS

**DPMS is just one part** of a general move toward protected mode DOS. Of course, for years there have been DOS extenders such as Phar Lap's 286 DOS Extender and 386 DOS Extender. DOS extenders continue to grow in importance—witness the incorporation of Phar Lap's 386 DOS Extender in most of Microsoft's new application software for DOS, including FoxPro, Microsoft C++, Visual C++, and Microsoft FORTRAN.

However, DPMS illustrates a somewhat newer trend: moving pieces of the operating system itself into protected mode. In addition to Novell, whose DPMS is likely to play an important role in Novell DOS 7.0 and NetWare 4.11, other companies are also working on moving TSRs and device drivers out of conventional memory and into protected mode.

In version 3 of its NetRoom memory management utility, Helix Software has introduced a "Cloaking API." Like DPMS, Cloaking is a method for moving TSRs and device drivers into protected mode. Helix has gone beyond this, providing cloaked versions of the system BIOS and video BIOS, both co-developed with Award Software. According to Helix, the cloaked BIOSs occupy only 8K of conventional memory instead of the normal 96K. The extra 88K becomes available as XMBs. The Cloaking API works as an extension to the EMS and XMS interfaces in real mode and uses INT 2Ch in protected mode.

Probably the most important effort to move pieces of DOS into protected mode is coming from Microsoft itself in Windows for Workgroups (WW). While WW may be relatively insignificant as networking software, it includes a number of critical device drivers (see Chapters 1 and 3) that require parts of MS-DOS. WW 3.11 includes VSHARE 386, VREDIR 386, and VIAT 386, which contains 32-bit protected mode code for SHARE, the network redirector, and even the FAT file system. These VxDs are also part of Microsoft's Chicago, which will provide an entire protected mode operating system.

## Novell NetWare

Novell's most significant products in the network software market are NetWare, a system NetWare, which includes both of the PC network software market: With MS Net, LAN Manager, and WW, NetWare has made three attempts to break into this market, with the Microsoft dominance in other PC software markets serving as a not entirely insignificant opening. As networks increase in importance, so will NetWare's success. Microsoft and Novell.

How does NetWare relate to DOS? NetWare, like servers running DOS, runs in real mode. The NetWare operating system, which was completely rewritten when NetWare 3.0 and higher supports multiprocessing, multitasking, and a form of dynamic linking called Loadable Modules (NLMs).

On a client workstation, however, the NetWare's client NLMs do not DOS. While DOS programs tries to open files, for example, NetWare detects whether the file is on a local or remote server rather than a file on the local disk. The access to NetWare files is directed to the NetWare file server. As we would see, NLMs make changes to the INT 21 interface, where it effectively replaces large parts of DOS, such as file classes, file buffers, and other DOS.

As Chapter 8 discusses in detail, DOS provides a standard interface (INT 21) to interface, or writing such software, called file system managers, the network redirector interface. However, NetWare 4.11 do start using the network redirector interface, until version 4.0.

## NETX and INT 21h

In NetWare 2.0 and 3.0, the workstation shell hooks INT 21h and looks for relevant calls. Nowell says this is this "shell" that can request. Because NETX hooks interrupts in front of BIOS, DOS, or Emule, a request NETX chains to the previous INT 21h handler, otherwise NETX sends a NetWare request packet to the appropriate file server. The NetWare shell communicates with multiple servers using the "proprietary" read undocumented NetWare Core Protocol (NCP). The workstation NCP packets to the file server using Novell's low-level IPX protocol.

NetWare's workbooks INT 13h, printing, INT 20h, error, error, and INT 21h. To keep track of disk operations from the old INT 20h and INT 27h termination and TSR routines.

When a file is opened, INT 21h AH 30h or AH 66h must be handled by a NetWare file server rather than locally by DOS. NETX converts the *drive mapping*. Any file open requests involving drive C might refer to the local hard disk, while a file open request for J: would refer to the file C:\JOBAR on a volume such as JSE:SYSPUBLIC on a file server. In addition, the workstation's drive mappings, NETX uses several tables: the Drive Flag Table, Drive Connect Table, and Drive Connection ID Table. Novell documents all these in the *NetWare System Architecture: A Guide to the User's Manual*, Chapter 5, Connection and Workstation Environment Services, Chapter 7, Directory Services.

Each set of files maintains that correspond to a workstation's 32 drives. These are drives A through Z, plus two additional reserved drives with drive letters and, unfortunately, Drew Major and the same, unexplained NetWare last apparently not read Dr. Nancy's workbook *On Beyond Z*. However, this source information on the others that come after Z. Whenever NETX's INT 21h receives a request for a file's directory request, it can consult the Drive Flag Table to see if the specified computer is local or remote. For remote drives, NETX can then use the Drive Connection ID Table to find the location into which the workstation drive is mapped. For local drives, NETX chains to the previous INT 21h handler.

When a DOS program opens a file on a NetWare server, it gets back a file handle, just as in DOS without NetWare. The file handle is really onto the File Table (FT) associated with the program's PNP. An undocumented Chapter 8 file tables are normally indexed into the System File Table. NetWare handles a file handle can be a negative number, starting with FFh and working backward. This treatment uses each entry into the internal file handle table that NETX maintains. In this case, the open file NetWare server is by the NETX, receive versions of NetWare, the reverse file handle, starting with FFh, can be specified in the CONSOLE.SYS. This means that the maximum file handle is 255 (0xFFh) minus the value of the file handle, a large value such as FFh-250 in CONSOLE.SYS. Like many other things, NETX would allow users to simulate DOS practice, this one goes to the NetWare 4.0 DOS Requests, where Novell uses the DOS redirector, thereby doing away with the need for separate file handles.

It is important to remember the goal here. A DOS program should be able to access files on a NetWare file server exactly the same DOS calls it uses to access local files; it should not need to be aware of the difference.

Of course, some applications will need to be NetWare aware. NetWare provides a large set of functions available to INT 21h. It is another reason why NetWare hooks INT 21h. For example, AX=NetWare File Open, Get Drive Handle, Table, AX=FF01h is Get Drive Flag Table, AX=FF02h is Get Drive Connection ID Table, and AX=FF04h is Get File Server Name Table.

Consider the ENUMDIR program in Chapter 8, which walks the DOS Current Directory Structure, displaying the contents of directory for each drive. Assuming that such a program is really using the file system, each local shell routine under NetWare, it would be important for the ENUMDIR to be aware of the CDS. A NetWare-aware version of ENUMDIR could walk each of the NETX tables to be sure that it now walks the CDS. It could also call INT 21h AX=FF20h Get Directory Path to let the directory handles into readable directory pathnames. Actually, while



ENU MDRV could use NetWare API calls to inspect the NETX tables; it would be easier to determine the names of redirected network drives and printers with DOS function INT 21h AX=5F02h Get Assign List Entry (see NETX DRV C in Chapter 8). This works because NETX handles the INT 21h AH=5Fh call.

Programming NetWare with its extended INT 21h calls is an enormous subject. For more information, see Charles Rose's *Programmer's Guide to NetWare* or the "Novell's Extended DOS Services" chapter in Barry Nance's *Network Programming* (see Ralph Davis, *NetWare Programming Guide* discusses the same subject, but using Novell's somewhat defective NetWare 3.12 interface library, rather than with direct INT 21h calls).

### NetWare 4.0 and the Network Redirector

NetWare 4.0 has a completely new architecture. As Novell notes, "technology outgrew the roughly seven-year-old NETX shell." Apparently, one reason Novell wanted to move to an INT 21h AH=11h redirector, and away from hooking INT 21h, is that it became necessary to modify NETX every time a new version of DOS came out. NETX is very dependent on the DOS version, whereas the many versions such as NETX, NET4, NET5, and so on, DOS 6.0 uses, by NETWR command, to tell NETX that it is supposedly running under DOS 5.0.

One of Novell's goals for NetWare 4.0 was to make the workstation client software more modular. Just as there are NLMs on the server side, NetWare 4.0 introduces Virtual Loadable Module (VLM) on the client side. A VLM is something like a dynamic-link library, as modules that can be loaded and unloaded depending upon which functionality a user needs. A VLM developer's kits available from Novell's developer services is a handy review.

NetWare 4.0 includes a NETVLM module, which provides back-end connectivity with the end front-end requests. However, this is only loaded optionally for applications that specifically require it. The required module is REDIRVLM, which is standard INT 21h DOS redirector.

Since this is a redirector, there isn't much more we need to say about it now (Chapter 8 explains this DOS standard in detail). When a front-end request is such as NETX objectives, it passes DOS and all the underlying DOS services such as the CDS and NLS, a complete sub-set part of DOS. Instead of hooking INT 21h, a new DOS redirector is called. DOS files on all drives, whether local or remote, are accessed through the CDS by NETX, regardless of how whether local or remote. Because a redirector doesn't bypass these DOS data structures, programs such as ENU MDRV could be used to connect to a NetWare 4.0 through a normal user, at least I want to know the NetWare side of the connecting between end and server.

On the other hand, NetWare's new workstation client is not a 100% redirector. Yes, it implements the INT 21h redirector, but we found some new NET 21h errors. NET VLM is only loaded. Some must call the hooking INT 21h, to support NetWare extensions such as the INT 21h AH=12h interface (see below). Novell's engineers refer to an INT 21h hooker as a "Shell" of course, to only that implements the INT 21h interface as a Redirector. The new DOS client includes both a Shell and a Redirector, so the file is "Novel DOS Redirector".

You might wonder how a redirector connects to NetWare, can it support the legacy NetWare drives through 27=52, since the DOS DOS uses FASTDRIVE. Yes, as the maximum block size. A Novel READVLM VLM file notes something like that "DR DOS could be enhanced to allow these drives to be treated as such by providing a new option for the FASTDRIVE parameter in CONFGSYS (for example, FASTDRIVE=37). Similarly, MS-DOS could be enhanced if Microsoft decides to adopt any changes made to DR DOS 7.01; otherwise, we could substitute the FASTDRV utility from Chapter 2 in this way.

On the other hand, perhaps the NetWare drives = Z just aren't that important. Novell has always intended these drives for internal use by programs that need to temporarily map volume and not generally for end users. Novell has long warned developers that the temporary file systems

roughly work in their place. Novell suggests using direct Universal Naming Convention (UNC) see Chapter 8, servers and file path aliases. For instance, rather than mapping `drive = *` to `\\SERVER\SYSMIBR`, a file access program can just use `\\SERVER` instead, can access them using `\\SERVER\SYSMIBR`. This functionality is available in NETX, though apparently with some problems in some DOS functions, notably Rename.

### How NETX Changes INT 21h

We can see that NetWare adds many handlers to the INT 21h interface. Novell documents some of these, and others are well documented but we know of them only through inspection of code. For example, some code in `\\SERVER\NetWare\Intercall\Intc`. Because the NetWare extension to INT 21h encompasses many topics we need instead look here at what changes NetWare makes to the standard INT 21h functions. For following discussion, please refer to information kindly provided by experienced NetWare users, especially of NETX, following the techniques discussed in greater detail in Chapter 6.

Of course, not all NETX modifications to INT 21h would disappear in the NetWare 4 release, because users would not even see it. In order to be installed, we had to hook INT 21h in order to support NetWare's own calls to INT 21h such as the `INT 21h` interface. The interesting question then is, "What does the standard NetWare 4 call address, and does some call take for the DOS world using NetWare's own file system support?"

NETX CALLS hooks INT 21h and puts its own code over that almost every single DOS function. This is a very interesting way to support NETX *de novo* changes, then to emulate at the times it changes. The following are the only INT 21h functions that NETX does not special case: 00h, 10h, 1D, 20h, 24h, 2Ah, 30h, 32h, 34h, 35h, 37, 38h, 49h, 54, 55h, 56h, 59h, 5D. NETX would try all other DOS functions, and since, for example, NetWare intercepts all I/O and file system calls, with those intended for NetWare file servers turned into SMP requests, as described above. This is a very interesting way to support NetWare's own calls to DOS, in a chapter on other DOSs. NETX replaces so much of DOS that by itself, outside the NetWare's own file system, running under a DOS, it is not a DOS. I address some of these changes in my INT 21h AM-140h to determine if they are really done by NetWare, but, since we measure such functions (PSP, INT 21h) using 4-4 below.

Of the changes, not all some of the less noticeable, serious changes that NETX makes to the DOS INT 21h interface:

- **Special Select of Handled Redirection:** For INT 21h functions 1, 2, 6, 7, 8, 9, 0Ah, 0Bh, 0Ch, 10h, 11h, 12h, 13h, 14h, 15h, 16h, 17h, 18h, 19h, 1A, 1B, 1C, 1D, 1E, 1F, 20h, 21h, 22h, 23h, 24h, 25h, 26h, 27h, 28h, 29h, 2Ah, 2Bh, 2Ch, 2Dh, 2Eh, 2Fh, 30h, 31h, 32h, 33h, 34h, 35h, 36h, 37h, 38h, 39h, 3Ah, 3Bh, 3Ch, 3Dh, 3Eh, 3Fh, 40h, 41h, 42h, 43h, 44h, 45h, 46h, 47h, 48h, 49h, 4Ah, 4Bh, 4Ch, 4Dh, 4Eh, 4Fh, 50h, 51h, 52h, 53h, 54h, 55h, 56h, 57h, 58h, 59h, 5Ah, 5Bh, 5Ch, 5Dh, 5Eh, 5Fh, 60h, 61h, 62h, 63h, 64h, 65h, 66h, 67h, 68h, 69h, 6Ah, 6Bh, 6Ch, 6Dh, 6Eh, 6Fh, 70h, 71h, 72h, 73h, 74h, 75h, 76h, 77h, 78h, 79h, 7Ah, 7Bh, 7Ch, 7Dh, 7Eh, 7Fh, 80h, 81h, 82h, 83h, 84h, 85h, 86h, 87h, 88h, 89h, 8Ah, 8Bh, 8Ch, 8Dh, 8Eh, 8Fh, 90h, 91h, 92h, 93h, 94h, 95h, 96h, 97h, 98h, 99h, 9Ah, 9Bh, 9Ch, 9Dh, 9Eh, 9Fh, A0h, A1h, A2h, A3h, A4h, A5h, A6h, A7h, A8h, A9h, AAh, ABh, ACh, ADh, AEh, AFh, B0h, B1h, B2h, B3h, B4h, B5h, B6h, B7h, B8h, B9h, BAh, BBh, BCh, BDh, BEh, BFh, C0h, C1h, C2h, C3h, C4h, C5h, C6h, C7h, C8h, C9h, CAh, CBh, CCh, CDh, CEh, CFh, D0h, D1h, D2h, D3h, D4h, D5h, D6h, D7h, D8h, D9h, DAh, DBh, DCh, DDh, DEh, DFh, E0h, E1h, E2h, E3h, E4h, E5h, E6h, E7h, E8h, E9h, EAh, EBh, ECh, EDh, EEh, EFh, F0h, F1h, F2h, F3h, F4h, F5h, F6h, F7h, F8h, F9h, FAh, FBh, FCh, FDh, FEh, FFh. This is a very interesting way to support NetWare's own calls to DOS, in a chapter on other DOSs. NETX replaces so much of DOS that by itself, outside the NetWare's own file system, running under a DOS, it is not a DOS. I address some of these changes in my INT 21h AM-140h to determine if they are really done by NetWare, but, since we measure such functions (PSP, INT 21h) using 4-4 below.
- **Redirection:** NetWare's own calls to DOS, in a chapter on other DOSs. NETX replaces so much of DOS that by itself, outside the NetWare's own file system, running under a DOS, it is not a DOS. I address some of these changes in my INT 21h AM-140h to determine if they are really done by NetWare, but, since we measure such functions (PSP, INT 21h) using 4-4 below.

### Figure 4-2: NETX Handler for INT 21h AM-0Eh (Select Disk)

```

4491 1006  SELDISK_0E
4491 1006  push es
4491:1007  cmp di,20h          ; drive == 32?
4491 100A  jae PASS_0E_TO_D05 ; yes, chain to previous INT 21h
4491 100C  mov bl,dl          ; no, put drive # into BX

```



with a PSP of 0FFFFh [1]. Running these versions of the C compiler under NetWare thus trashes not only the INT 0Bh vector, but also some other location that depends on whether the A20 line is enabled. If A20 is off, 0FFF 003E wraps around to 0000 002C, which holds the INT 0Bh vector. COM2 uses INT 0Bh. Imagine someone's surprise when, after compiling a program with QuickC under NetWare, their communications package fails! The DOS world, as a constant reminder of how "too many cooks spoil the soup," is full of interesting interactions like this. It almost makes you wish that Microsoft had even more of a monopoly.

Several friend rumors that QuickC performed such a test, but found it so difficult to believe we had to check it out for ourselves. Digging out a copy of version 2.51 of the Microsoft C compiler with QuickAssembler, we found that the QuickC README.DOC file actually refers, in passing, to this test, saying that the product "prints a warning message when a non-Intel version of DOS is detected. The /noingn option prevents this message from being printed." Indeed, QC.EXE (dated April 6, 1990) contains the following message, which sounds strikingly reminiscent of some of the unfair trade practices explicitly banned by the Clayton antitrust act:

```
WARNING: This Microsoft product has been tested and certified for use
only with the MS-DOS and PC-DOS operating systems. Your use of this
product with another operating system may void the available warranty
protection provided by Microsoft on this product.
```

Remember, this message is coming from a Microsoft C compiler, not from part of MS-DOS. It is no wonder that shortly after bringing this warranty-related warning message to the attention of an attorney at the U.S. Federal Trade Commission (FTC), one of the authors received a letter informs that the FTC's consumer affairs bureau had suddenly become involved in the investigation of Microsoft, and that the Magnuson-Moss Warranty Act had entered the picture.

What is a good thing that this message can be turned off, what about those calls to SetPSP() and SetSPN()? Recall that, when this code was written, the SetPSP function was fully documented, yet Microsoft's C compilers compete with products from other vendors. It is not clear whether SetPSP() defines a certain fixed spot in the DOS data segment, a just the standard interface test case of insider knowledge that Microsoft applications are not supposed to be taking advantage of, particularly when the sole purpose of this code is to make Microsoft users think they shouldn't run versions of DOS other than Microsoft's. You could make a good case that it is dumb to bother with anything except genuine MS-DOS, but it is appropriate, and possibly even restraint of trade, for Microsoft's C products to try to convince you of this point.

To view QC.EXE's early inclusion of such code, we first ran it under the following INTRSPY script:

```
; psp0.scr
intercept 21h
function 50h on_entry
if ((ba == 0) || (ba == 0FFFFh))
output SET PSP = %b from %C %;" IP
```

Sure enough, INTRSPY produced the following output, indicating that the current PSP was being set to 0 and 1 (FFFFh):

```
SET PSP 0 from 9E33 016F
SET PSP FFFFh from 9E33 0177
```

Using the CS:BP addresses the NTRSPY script displays, we used WINICE to disassemble the code. Figure 4-3 shows the results.

**Figure 4-3: Microsoft QuickC DOS Detection Code**

```

u 89e33.014e
; this is in the middle of some sort of "is dos okay?" function
9E33.000014E MOV AH,52
9E33.0000150 INT 21 ; get DOS DS into ES; ignore BX
9E33.0000152 MOV AX,ES-[0004] ; DOS_DSE[4] = SDA type
9E33.0000156 CMP AX,0001
9E33.0000159 JA 018C ; hmm, DOS DSE[4] * 1 is ok, DOS 2?
9E33.000015B SHL AX,1
9E33.000015D MOV SI,AX ; AX is 0 or 2
9E33.000015F MOV SI,[SI+0134] ; GC knows CURR_PSP ofs in DOS_DS
9E33.0000163 MOV AH,51 ; [134h] + 020e, [136h] 0330h
9E33.0000165 INT 21
9E33.0000167 MOV DX,BX ; save current PSP in DX
9E33.0000169 XOR BX,BX ; BX = PSP = 0
9E33.000016B MOV AH,50
9E33.000016D INT 21 ; SetPSP(0)
9E33.000016F CMP ES[SI],BX ; see if DOS_DS[Curr_PSP] matches
9E33.0000172 JNZ 0183 ; no not MS-DOS
9E33.0000174 DEC BX ; BX = PSP = 1
9E33.0000175 INT 21 ; SetPSP(1)
9E33.0000177 CMP ES[SI],BX ; see if DOS_DS[Curr_PSP] matches
9E33.000017A JNZ 0183 ; not not MS-DOS
9E33.000017C MOV BX,DX
9E33.000017E INT 21 ; restore PSP
9E33.0000180 JMP 018C
9E33.0000182 NOP
; not_msdos
9E33.0000183 MOV BX,BX ; problem, no match
9E33.0000185 INT 21
9E33.0000187 XOR AX,AX ; return false = 0
9E33.0000189 JMP 0184
9E33.000018B NOP
; is_msdos
9E33.000018C XOR AX,AX ; ok match
9E33.000018E DEC AX ; return true = -1
; done:
9E33.000018F POP ES
; ***
dw 89e33:0134 0136
9E33.0000134 020E 0330

```

This code definitely is taking advantage of inside knowledge, and is using a narrow definition of DOS compatibility. For example, QuickC wants to find the address of the current psp field in the DOS data segment. While also and ironically, the common way to get this in third-party code would be to call INT 21h AX=5006h to get a pointer to the Swappable Data Area and then look at offset 10h. Instead, the code calls INT 21h AH=52h to get the DOS data segment into ES (ignoring the SysVars offset in BX), then uses a little-known value at offset 4 in the DOS data segment to determine whether there is a DOS 3.0-style SDA or a DOS 4.0-style SDA. Based on this, the program uses hard-wired offsets for the current PSP in the DOS data segment (offset 020Eh in DOS 3.0 and offset 0330h in DOS 4.0 and higher).

Despite the supposed separation of church and state between application programmers and operating systems programmers at Microsoft, the use of DOS\_DS[4] reveals an intimate knowledge of the DOS code. Unlike many other features of undocumented

DOS which are open secrets, the SDA type indicator at DOS\_DS[4] has not been publicized at all. To our knowledge, the only place this has been discussed in print has been in the writings of Geoff Chappell (see the discussion of "Direct Access to Kernel Data" in Chapter 13 of his *DOS Internals*). The code for SHARE and other DOS utilities relies on DOS\_DS[4], but it is certainly odd to see it showing up in the code for QuickC, which is an application, and doubly odd to see it showing up in the code for the nasty QuickC code. Of course, PSPTEST first checks to see if NETX is running, if it is, PSPTEST asks the user if they really want to run the test, since SetPSP(0) under NETX may corrupt their system.

Listing 4-4 shows a short C program, PSPTEST.C that duplicates the nasty QuickC code. Of course, PSPTEST first checks to see if NETX is running, if it is, PSPTEST asks the user if they really want to run the test, since SetPSP(0) under NETX may corrupt their system.

#### Listing 4-4: PSPTEST.C

```

/*
PSPTEST.C
Andrew Schulman, June 1993
From Undocumented DOS, 2nd edition (Addison-Wesley, 1993)
*/
#include <stdio.h>
#include <stdlib.h>
#include <ctype.h>
#include <conio.h>
#include <dos.h>

#define PK_FP
#define PK_FP(seg, ofs) \
    ((void far *) (((unsigned long) (seg) << 16) | (ofs)))
#define I

typedef int BOOL;
typedef unsigned char BYTE;
typedef unsigned short WORD;

#pragma warn -rvl
WORD _dos_getpsp(void)
{
    asm mov ah, 51h
    asm int 21h
    asm mov ax, bx
}

void _dos_setpsp(WORD psp)
{
    asm mov bx, word ptr psp
    asm mov ah, 50h
    asm int 21h
}

#pragma warn -rvl
WORD _dos_getds(void)
{
    asm mov ah, 52h
    asm int 21h
    asm mov ax, ax
    // ignore SysVers offset in BX
}

// Novell Return Shell Version function (INT 21h AH=[AH AL=01h])
// see INTRLIST on accompanying disk, also see Barry Nance,
// Networking Programming in C, pp. 117, 341-2.
BOOL network_shell(void)
{
    char buf[403];
    char far *fp = buf;
    asm push di
    asm mov ax, DEAD1h
    asm mov bx, 0
}

```

```

_asm les di, fp
_asm int 21h
_asm xor ax, ax
_asm mov ax, bx
// If BX still 0, then NetWare not present, return in AX
_asm pop di
}

void fail(const char *s) { puts(s); exit(1); }

main()
{
  BYTE far *dos_ds,
  WORD far *curr_psp_ptr,
  WORD psp,

  dos_ds = (BYTE far *) HK_FP(_dos_getds(), 0),
  switch (dos_ds[4])
  {
    case 0  curr_psp_ptr = (WORD far *) (dos_ds + 0x020E), break,
    case 1  curr_psp_ptr = (WORD far *) (dos_ds + 0x0330), break,
    default: puts("DOS_DS[4] > 1"); goto okay;
  }

  if (network_shell())
  {
    puts("Under NETX, SetPSP(0) can corrupt your system ");
    fputs("Still want to continue? [Y/N] ", stdout);
    if (toupper(getch()) != 'Y') fail("bye"),
    putchar('\n');
  }

  psp = _dos_getpsp(),
  if (*curr_psp_ptr != psp)
    fail("curr_PSP field wrong!");

  _dos_setpsp(0),
  if (*curr_psp_ptr != 0)
  {
    _dos_setpsp(0),
    fail("SetPSP(0) test failed");
  }

  _dos_setpsp(-1),
  if (*curr_psp_ptr != -1)
  {
    _dos_setpsp(-1),
    fail("SetPSP(-1) test failed");
  }

  _dos_setpsp(0),
  puts("SetPSP(0) and SetPSP(-1) tests succeeded");
okay:
  return 0;
}

```

It's worth noting that this program succeeds, not only under MS-DOS 3.0 and higher, but also under DR-DOS 5.0 and higher. In other words, DR-DOS keeps the current PSP at exactly the same location, DOS\_DS[02DEh], as DOS 3.0 does. What a coincidence!

PSPTEST also succeeds under the old Concurrent DOS 3.86 2.01 which doesn't implement the DOS\_DS[4] flag. QuickC also runs in this environment. If you examine the code, you will notice that any environment with DOS\_DS[4] > 1 is okay, presumably this is because QuickC has to run on DOS 2.x, which didn't implement the DOS\_DS[4] flag. In

fact it's unclear under what circumstances you would have a DOS clone in which (DOS DS[4] = 0 but DOS DS[2D6h] != CURR PSP) or (DOS DS[4] == 1 but DOS DS[330h] != CURR PSP).

In summary, the SetPSP() code seems to package up all the problems of undocumented DOS + a nutshell we have NetWare hijacking unused fields in the PSP (or perhaps not—see below), NetWare failing to think through all the implications of hooking INT 21h, Microsoft applications accessing undocumented DOS kernel data, and Microsoft applications testing for the presence of non-Microsoft operating systems in what certainly looks like a messy timing arrangement (see Chapter 1).

What about that Novel task ID? According to an engineer at Novell, NetWare uses its ID to keep track of multiple programs running on a single workstation. Just to make things interesting, this is not the PSP's unique identifier reserved for Novell's *NetWare*®; Microsoft's source file for the PSP ID is found in the OEM Adaptation Kit (OAK; see Chapter 2) and contains a "PSP Novell Use C" field and a "Novell Use C" field. Novell's *NetWare* team uses these. The task ID starts at 1 for the master copy of COMMAND.COM and increments for each task loaded that does DOS I/O. Some programs are associated a task ID by NETX.

- *IRUFNAME* (INT 21h, vector 00B): NETX hooks the IRUFNAME function, and uses it to determine a NetWare server's partition name. Thus, you can call IRUFNAME to get the working server and volume name for a NetWare drive.

As a result, these NETX modifications to INT 21h basically present us with a different version of DOS. On the other hand, these modifications go away starting in NetWare 4.0.

### Undocumented NetWare

NetWare has vast amounts of undocumented functionality—enough to fill a coffee book. As noted earlier, *NetWare: A Guide to the Undocumented* is scheduled for publication in 1994. In the meantime, it is interesting to consider the following:

- *NetWare*™: *NetWare*™ (the name of the system used on NetWare file servers) is proprietary, which makes it difficult to write products such as Norton Utilities or PC Tools for NetWare file servers. Any company vendor (Checkmate has recently acquired the NetWare file system).
- *File servers*: Many of the *NetWare*™ file servers, the names of which appear in the Watson linker tool, Novell steps. Many of these servers come in flavors with an EE prefix (indicating "Enhanced") and could be used either for writing alternate OS systems or for utility vendors interested in continuous backup, server backup software, and the like.
- *NetWare File*: The API for NetWare File appears to be completely undocumented. The Inter-File File System (comparing task 1) is a low NetWare File functions, such as INT 21h AX 1000h (CL), INT 21h AX 1001h (CL), and AX 1002h (SERVIAL) (not a check). This may all change with Version NetWare™, which is scheduled for the next release of NetWare File. The server component will be based on the same TCP protocol as NetWare, but on the special File protocol, which is not used and work with the NetWare File common client for DOS being shipped with NetWare 4.0. Clients can start protocols, including both NETX to be a NetWare client and CLIENTX to be a File client. With the common client you will be able to use one IPX instead of two. Rumor has it that Novell DOS 7.0 will virtually incorporate Personal NetWare for DOS and that the peer-to-peer server will use DPMS.
- *The NetWare Core Program*: Novell considers NCP "proprietary." While NCP is immediately usable through a proxy with a network sniffer such as General Software's The Sniffer, or even



Now I own LANalyzer. Novell has never publicly documented the protocol. A partial description appears in *Novell's Guide to NetWare 3.11 Analysis* by Laura Chappel, Scribe, Novell Press, 1993. General Software's *The Networker* comes with IPX socket code, including an NCP module. For more information on NCP, see Peter Neumann's article "Novell's Unclim Hands: The NetWare Core Protocol NCP" in *The Info Journal*, November 1993.

- The I2 interface.** A novel feature, NETV extends INL 21h to provide many NetWare-specific functions. While Novell documents many of the extensions, some (including INL 21h and I2h) are not documented. It also sends programs to scan raw NCPs (with the filename I2net) in order to issue raw NCPs as the *IPX Drivers / Drivers* article cited above notes you can do directly via IPX. A friend of a Microsoft complains talk about the potential for the "leaky black" "This one must already been documented, although they have released a few code fragments that actually use the interface. This was a response to developers who were explicitly denied that they could not do something that some Novell utility could do. However, Novell does seem committed to supporting I2 and has provided attendees of its "BrainShare" developer's conference with information on the interface, though at the same time keeping NCP information without which I2 is useless under wraps. An engineer at Novell reported the following information on the I2 interface:

```
RawNRequest      ;242 (F2)
** undocumented **
Generic function used to send raw NCP requests. Documented in part
by the Erase Files function for purging files (INT 21h A3+F246h)
input:
AL      NCP function request code
DS:SI   Request List
CX      Request Length (in bytes)
ES:DI   Reply List
DX      Reply Length (in bytes)
output:
AL      Return code
```

Other undocumented INL 21h calls include AttachEject, AH-B4h, an undocumented Novell provides to give DOS the access to NetWare files, including on Macintosh OS 7, and INL 21h GetData (AH-B5h) "a nice little jewel" for the UNIXWARE 386 Windows VxD, with simultaneous AH-B6h to get instance data, start and exit critical machines, and so on, and ReturnSearchContext (AH-B7h) used to exit the master environment "4 COMSPEC" and "PATH" parameters.

The trickiest one in the master environment is interesting given the difficulty noted in Chapter 10 of hiding the master environment. According to a engineer at Novell, they use some directly out of the Spontaneous Assembly library, as they also seems to work best. Novell changes the "PATH" when you log in to enable cross-search drives in the set. Obviously, to make this change, Novell must be in master environment. According to one user, this makes for some real nastiness if you try to LOGON to Novell while shielded out of another program. So now Novell is providing a undocumented way that other programs at NetWare find the master environment back for them.

## OS/2 2.x: "A Better DOS Than DOS"?

When OS/2 1.0 first appeared in 1987-88, Microsoft and IBM, along with many PC software developers, thought OS/2 would replace MS-DOS and take over the desktop. OS/2 1.x had a rich, robust DOS compatibility box, commonly referred to as the DOS cell or priority box, or C-memo, in which ran some DOS software and on which much work was spent. But it's indicative of a system



vendors." Presumably users of V86 mode, such as Microsoft, IBM, Novell, Quilias and Quarterdeck, will have no trouble getting their hands on the necessary documentation from Intel and incorporating these Pentium enhancements into their software. Then it will just be a question of when end users actually get machines with Pentium processors!

There are many interesting and important aspects to MVDMS, OS 2.2.0. Here, however, we focus exclusively on how well OS 2 emulates MS-DOS, emphasizing, of course, its support for undocumented DOS functions and data structures. For more information on MVDMS technology, see the chapter on compatibility of Harvey Deitel and Michael Kogan's *IT: Design of OS 2* (Addison Wesley, 1992).

### MVDMS and VDDs

OS 2 MVDMSs wear a first a lot like DOS boxes in Windows Enhanced mode. In both environments, DOS programs make normal INT 21h calls and INT 2Bh calls do INSDirect calls to ports, read and write memory locations, and so forth. Other software stuff (i.e., DOS programs) looks, to a virtual machine monitor, exactly like these events and handles them out to various device drivers. VDDs in Windows, but atomized as VDDs in both OS 2 and Windows NT, which is the only software device the DOS program thinks it's direct "talking to."

These devices aren't necessarily pieces of hardware. For example, just as Windows has the V86MMIO VDD to emulate the XMS and EMS specifications that's right, XMS calls in Windows aren't handled by HIMEM.SYS, V86MMIO, or real-mode XMS, as in OS 2's design, but HIMEM.SYS or an FMM driver, instead. To fix the VDDs XMS.SYS and XMM.SYS. In OS 2, VDDMM.SYS provides the DOS enhanced Mode Interrupts. VDI.XMS provides optional protected mode IN, 21h service, similar to DOSMMIO.Windows.

Of course, some VDDs in OS 2 are some VDDs in Windows. For example, there are VPI.SYS similar to VPI.D in Windows, VAMEM.SYS, VDD in Windows, VMOEM.SYS, VMD, VKLD.SYS, VKD, and so on. Here, and many sometimes in Windows Enhanced mode, you get down to a hardware or some virtual device or device in Windows. VDDs use the undocumented Linux Executable File table of the OS 2 VDDs as the EXE table that, which is a hacked version of a FAT file system, is sparsely documented (see FAT32.H in the OS 2 2.0 SDK CD-ROM). VDDs in OS 2.2.0 do VDDs in Windows that offer, but not real mode, the NT VDDs which we see in a real mode 3.0 mode.

VDDs apply a bit to support OS 2's system boot, a user-visible, through VDDs in the same way that VDDs can provide any type of user-visible programs running under Windows. To running software interface a VDD to OS 2, you get only the VDD's DOS program, just as VDDs in Windows can call services from VMM or other VDDs. OS 2 VDDs have a their special set of Virtual Device Helper VDH functions. For example, VDDs can hook V86 interrupts with VDHInstalIntrHook, interrupt hooks, and other services with VDHInstalPMIOHook. A DOS program's I/O port access can be controlled with VDHInstalIOHook. The OS 2 2.0 *Virtual Device Driver Reference* (IBM work) on the High 386, documents these VDH functions, as does the VDD chapter of Steven C. Morrison's work *Windows OS 2.2: A Design Document*.

All this is very similar to VDD programs in a Windows Enhanced mode. This isn't surprising since Microsoft wrote the OS 2 in this mode, so the IBM-Microsoft device. When Microsoft trashes OS 2 in public, it's most likely to be in your real mode.

But there are some serious differences between OS 2 MVDMS and Windows 4.0 in real mode. We saw in Chapters 1 and 3 that Windows reflects many DOS calls down to MS-DOS in V86 mode. But in OS 2, DOS calls are not. This is a rather obvious point, but it bears explaining. OS 2 does not run on top of DOS; DOS is nowhere in sight. Any INT 21h calls that a DOS program makes under OS 2 must be handled by OS 2 itself. Windows, at least without V86MMIO, can't do this.

IO calls, switching a direct API translation and then calling down to DOS in x86 mode, but OS 2 does not use DOS file system calls internally to the OS 2 kernel. For example, INT 21h 31h 31h calls must be handled by the OS 2 DosOpen() function. INT 21h 31h 31h must be handled by DosRead() and so on.

OS 2 DOS programs running under OS 2 really do think they are running under a version of DOS. I will discuss all the documented INT 21h and INT 21h functions and many of the undocumented functions in a separate article on INT 21h API 30h and get back to DOS version numbers.

### So What Version of DOS Is This DOS Emulation Pretending To Be?

OS 2 2.00 [NT, 386] returns 0014h (DOS 2.00) on OS 2 2.10. It returns 0A14h (DOS 2.10) on OS 2 2.10. This is a common way for OS 2 to identify the DOS box. OS 2 1.1 is probably the first to be DOS 1.0. The original thinking here was probably that there was *no* real generic DOS 4.0x or 5.0x compatible software. Only MS-DOS 6.0 seems to have a version number that increases, and even then it may not be possible to run MS-DOS 6.0 on the hardware it seems possible that one day you will see a DOS 10.0 (or 20.0) at least. Does seem to be a reasonable value. As for one of our competitors, As of press time, a generic MS-DOS 10.0 (look at SERVER) can only handle INT 21h 31h 31h 09h. Thus, MS-DOS has no place to ever reach DOS 10!

The DOS version number is useful for telling DOS applications they are in fact running under OS 2 2.20 or under MS-DOS 6.0 or 6.44. 2.20 is not a real DOS version (from INT 21h AX 3306h) since MS-DOS 5.00 seems to not have a real version (it returns the same version as INT 21h AX 3306h) but the presence of function 6306h, which came about with DOS 5.0, shows that OS 2 is emulating the MS-DOS 5.0 or higher. Unfortunately, there is no call to get the simulated DOS version. An engineer at IBM reports that

The OS 2 2.0 DOS box was originally a copy of DOS 3.3 modified to only removing the OS 2.20's support for emulating DOS 5.0. We looked at the function specifications for INT 21h and DOS 3.3 and still work very well, except to be sure. Most dos calls are almost identical. The biggest one is to make sure the IO controller ID is a simple copy of the io controller ID of the disk. When we finished, we noticed that there was a DOS 5.00 version (from INT 21h 31h 31h 09h) and AX 4406h and AX 6411h) that were not in the original DOS 3.3 documentation, but publically as a DOS 5.0 compatible.

OS 2 MS-DOS 5.0 and higher is OS 2 2.0 VDM compatible like DOS version numbers. OS 2 2.00 is not compatible with the DOS VERSION setting. With DOS VERSION, you can set the DOS version to be DOS 2.10 or be told the version running under DOS 5.0 or a lower DOS VERSION. In the SERVER, MS-DOS accepts the DOS VERSION includes the version of the program that is running on version number set as the number of tracks to take the hard disk number (change disk size). As for the case in the version of MS-DOS 5.0 also has the same the number. As for the SERVER, changing the DOS VERSION field does nothing since OS 2 does not use it. It has a version of DOS. It simply determines what call of INT 21h API 31h 31h 31h 09h the specific application.

DOS VERSION is not one of our settings in the OS 2 VDMs. For example, VDIS ASIDEIVE is not one of our settings, drives in a VDM 2.0 does not use a Car read function. Most of our settings in DOS 5.0 are the OS 2 kernel handles all INT 21h 0x 0x 0x and does not use the settings. This setting for example, the same purpose as C:\DRIVE1 as a DOS CONFIG.SYS setting on other operating systems. Different VDMs can run simultaneously, each with a different DOS LASHDRIVE.

DOS 1.11 is not a real DOS version, similar to the FILES statement in DOS CONFIG.SYS. As discussed in Chapter 8, MS-DOS FILES sets the size of the system file table (SFT), a file handle in MS-DOS is a device of the program's job in a table (BT) which in turn only index into the SFT. In OS 2 VDMs don't use SFTs. While a file handle is still an index into the

JEI. On JEI boards OS 2.5 handles root SEI indices. From a DOS program's perspective, the SEI under OS 2.5 can't be much greater than 700 entries. Thus changing the SEI size is meaningless in an OS 2 VDM. While DOS FILES does that, it also increases the size of the VDMs. In a real sense, OS 2 DOS/MAINT functions, which are similar to INT 21's MAINT, is DOS.

Another example is DPMI-DOS API, which controls whether the DOS Protected Mode Interface provides the MVDM simulated mode DOS, a shared protected mode. As we saw in Chapter 3, DPMI and protected mode INT 21 have separate services. WIN OS 2 copes with DPMI and DPMI-DOS API with appropriate ones, which is Command C. It also interprets the presence of DPMI to mean that protected mode INT 21 is automatically available.

There's something special about these settings: they are simply registers. In a protected VDM, you document VDI registers (see Registering OS 2 DOS Settings below). For example, VDPX SYN registers DPMI-DOS V.2, VLDOS SYN registers DOS FASTDRIVE, and DOS FILES and DOSKRN<sub>x</sub> registers DOS VERSION. In a sense, these are just like SYSTEMINI settings in Windows. Excepted mode VDMs are also free to establish

## Registering OS/2 DOS Settings

There is a VDD call, VDHRegisterProperty, it's documented in the OS 2 DDK, to set up a DOS Setting. Any VDD, whether system supplied or user written, can register a property. One of the pieces of information you must supply is whether the property can be changed at any time or only before the VDM is created. DOS Settings can be viewed as similar to DOS CONFIG.SYS processing, except that in DOS, unlike with OS 2 VDMs, none of the settings is automatically changed once the machine is up and running. (We've seen in other chapters, of course, how to modify FASTDRIVE, by changing Sys Vars, or how to modify FILES, by adding new files to the SET file.)

The following minimal VDD demonstrates a user-written DOS Setting. It doesn't do anything but add GEORGE\_FULK to your notebook for DOS settings. The two reserved fields are really case sensitive: help is definitely the correct PMVDM.PDI, which runs the notebook, ignores them. The reader will not be surprised to learn that this code was supplied by George Fulk, a engineer at IBM who contributed much of the information on OS/2 VDMs in this section. (Software also of IBM supplied additional information.)

```
; GEO_FULK.ASM
386P
;-----
ASSUME CS CSEG, DS FLAT, SS FLAT, ES FLAT
CSEG SEGMENT DWORD USE32 PUBLIC "CODE"
Validate proc near
    xor     eax, eax
    ret
Validate endp
CSEG ENDS
;-----
ASSUME CS:CINIT_TEXT
DSEG SEGMENT DWORD USE32 PUBLIC "DATA"
PROPERTYNAME db     GEORGE_FULK , 0
DSEG ENDS
;-----
EXTRN VDHREGISTERPROPERTY: NEAR
CINIT_TEXT SEGMENT DWORD USE32 PUBLIC "CODE"
_VDDInit proc near
    push   ebp
    mov    ebp, esp
```

```

--- register the DOS setting
push  offset FLAT:PROPERTNAME
push  0 ;reserved
push  0 ;reserved
push  0 ;boolean
push  0 ;"other", non-system VDD
push  1 ;only change before VDM creation
push  0 ;default value, off
push  0 ;no range checking for boolean
push  offset FLAT:Validate
call  VDMREGISTERPROPERTY
or    eax,eax ;if VDM called fails,
; then fall the VDD init.
je    short InitDone
dec   eax ;init!=0 means success
;init==0 means failure
InitDone
    eave
    ret
VDDinit endp

[INIT text] ENDS
END _VDDinit

```

### Loading a Genuine DOS

As a alternative to using DOS emulation OS 2.2 you can also load an actual copy of MS-DOS. The IBM *Loadable File and Exec Capabilities: Running Specific DOS from Within OS 2.2.0* which was designed to deal with these systems, such as DOS EAX Requests, that must run a flat virtual machine copy of DOS 4.0 (also than with OS 2 DOS emulation), the procedure involves running a DOS boot block from within a diskette image (MIB) file with the OS 2 VM, VDK object code that can load the OS 2 USER SYSTEMS dev, which allows the specific copy of DOS running in DOS box to access the OS 2 file system.

So, I did some DOS calls to show that looks like 21h. How then does it communicate with the OS 2 kernel? Based on 6 of USER EXEC also some calls to INI 06h with ANSA as I said, the only thing is VM COM VDK as I did. However, this doesn't seem sufficient for a DOS kernel to do any thing. So, I've seen a VM to communicate with OS2KRNL running at Ring 0 to protect memory. USER EXEC uses HI instructions to communicate with OS2KRNL. We asked one of the IBM MVDL engineers how this works.

HI is our magic pin. The HI instruction is a Ring 0 instruction since V86 mode is Ring 3 and it works. A part of OS2KRNL called EXM86 Emulation 8086 receives the HI and takes action. EXM86 looks at the request and the two bytes following the HI to determine what the request means and not vice HI in its reaction. The code uses MIB to get the address of the VM object from the instruction to call to OS2KRNL, and get the MVDL part of OS2KRNL, which is the valid address.

Most requests to 8086 provide the VM Ring 0 and VM86 emulates the 8086 software using protected instructions. It's special operation HI is used MVDL also part of OS2KRNL is used to support the signaling and action. There are second magic pin EXM86 uses the VM object code VDD stubs such as RPL DOSKRNL and USER EXEC HI.

More similarities to Windows Enhanced mode uses RPL for the same reasons as the EM86 VM Exec Protected Emulation. But, we did see that DOS device drivers running under Windows NT, which are a cross between DOS and Windows, do not communicate with NT VxDs.

Using the OS 2 user EXEC to load actual DOS in V86 mode emulating DOS is tricky. The user EXEC is not a actual copy of DOS makes it more difficult to control pushing in the DOS kernel. In all states cases, I addition, USER EXEC must catch DOS calls at a lower level than the normal DOS emulation code.

### OS/2 2.x and Undocumented DOS

IBM's OS/2 *Compatibilities International* booklet, while noting that many DOS and Windows programs run under OS/2 2.0 modes, however, that "some applications may still become either design or because they use real-mode interfaces will not operate properly with OS/2 2.1.

As examples of DOS software that won't run under "actual DOS" the DOS™ IBM gives us to list "application programs that directly address the physical sectors of OS/2—managed non-removable DASD"—in other words, the Norton Utilities' timing-sensitive DOS programs. VCI clearly OS/2 does support the VMS LMS and DPMI specifications. DOS debuggers that set 386 hardware breakpoints on 386 DR0-DR7 register access, and Windows programs that rely on VxDs (as it includes an Win32s application since OS/2 can't load Win32s 386).

It's worth noting that more and more Windows applications require a VxD or Win32s. For example, Microsoft Visual C++ comes with three VxDs: CAVI386, DPMAN1386, and MMIO386. PC Week (April 26, 1993) had an extremely confused discussion of OS/2's ability to run VC++, indicating that VC++'s need for VxDs was somehow "excusable." At times at the IIC appeared at one point to interpret this to mean that VC++ perhaps required the VxDs simply to keep VC++ from running under OS/2. There is no evidence for this at all. CAVI386 is Novell's VxD for running the CodeView debugger on a single screen, and DPMAN1386 is set by the Par Tap DOS extender used by the VC++ command-line tools. If you're Windows applications won't run under OS/2 because the application requires a VxD, then it seems to be IBM's responsibility either to write a VxD (either for WIN OS/2 or to develop OS/2 VxD equivalents for the most important Windows VxDs).

OS/2 2.1 has improved support for DOS and Windows. You can launch DOS programs from WIN OS/2, and Windows Enhanced mode is supported to the extent of carrying WinMMS32 applications. Still no VxDs, however. Some (incidentally) that programs that require VxDs won't run under Windows NT either.

On the question of timing-sensitive applications, one of the MVDI developers at IBM writes:

As well as used DOS applications, we had to consider what would happen if a DOS application tried to use MVDI. The only exception was applications that requires hardware capabilities beyond our capacity. For instance, we estimate every gig including I/O ports. As yet, we can only handle 1000 interrupts/second. Some applications want a total of 6 megs. The I/O ports are produced in a 1000/60000 interrupts/sec. What if we could not we could never begin that rate. One suggestion I've received is to add in priority to DOS setting to give each sub-control of a piece of hardware to a session, or a second I/O ports. No other session in the system could use those I/O ports. For we could increase the interrupt rate dramatically.

It's important to know what kind of hardware support DOS emulation provides, or the whole reason for these elaborate virtualization schemes is that popular PC applications bypass approved methods, making low-level calls instead of high-level ones, accessing I/O ports and memory locations instead of making DOS calls, etc. This is necessary in the DOS world. The performance boost from bypassing DOS is the reason why these programs became popular in the first place. However, it also puts a great burden on the developers of environments such as OS/2, Windows NT, and Windows enhanced mode.

If as optimizations played with their efficiency, the world be case for DOS emulations. You would need all this elaborate device virtualization in the first place. Operating systems such as OS/2 could trap INTR21h and have a 31h handler, inside the INTR21h handler, everything could be done in a 32-bit mode. Virtualization, since impressive, hooking an I/O port just like it was an "intercepting" can lead to lower performance.

cover that use of low-level code is a fact of life in DOS software. Let's focus on what support OS 2 provides for undocumented DOS. The following useful topographic summary comes from an engineer at IBM:

We are supposed to support all documented INT 20h-2Fh calls, and most of the undocumented ones too. We are close, but not quite there. There are a couple of undocumented calls we can't make do with. The undocumented ones are, as a wild guess, 50% here.

Network redirectors, not yet.

All internal data structures are the same as DOS 5.0, but some are read-only and others are write-friendly placeholders. That doesn't mean that they are always used, just they have the structure and appearance. Many of the writes are read-only.

CDR or SRTIME, except for the INT 2Fh AX=1716h id, which we take. Of course, no OS 2 file system can do CDR, but we do not want to mess with a DOS boy.

Call to the MSCDEX. Nope, MSCDEX is a redirect. Note, however, that the ACROMSYSAD (the CD-ROMS) does support the MSCDEX INT 2Fh calls.

INTs work just fine. No knowing problems or restrictions.

The basic DOS calls and files should work out of skates, but not on the hard disk. No absolute addresses. INT 26h is permitted to the hard disk. Unless an application knows it is not to do so, most addressing operating system can pass it this.

In short, the OS 2 VDMs provide a read-only list of lists, SevVars, and SDA, but no CDS or MIB are to work with. Reusing some of the programs from other parts of this book under OS 2 2.0 confirms this point:

- OKAY: Chapter 2 works, so INT 2Fh AH=52h is supported, and SevVars has a valid LASHDRIVE value in the correct location.
- CDROM: Chapter 3 works, so SevVars also contains a pointer to a valid MIB chain. Of course, this MIB chain represents only one example of VDM memory, not the entire OS 2 address space.
- DEV: Chapter 7 works, so SevVars contains a valid SRTIME header pointing to a valid device in its system. This is only for DOS device driver chain, or for VDM including device headers such as VDDIR and DevDir, so, since OS 2 places a device driver VDDIR into a separate file, a valid header for DOS program may not be absent. (Warning: In OS 2 Device Chain DevDir, however, August 1990 issue, is a separate VDD chain.)
- EFILES: Chapter 8 fails, either because the SRTIME is absent in the VDM. Note, however, that OS 2 supports the use of SRTIME and MIB files. Labels MIB files, sources which are used by OS 2 applications such as DevDir and DevDir. (See *DevDir* in *OS 2*, p. 236.)
- FN: MIBV: Chapter 8 reports, can't get CDS\*? Again, OS 2 itself does have a CDS string, but it is not accessible as a program's basis in the Protected Data Area (PDA), see Acrol and Kogan, p. 236.
- PDA: COM: Chapter 8 fails, on average, but the network redirector interface isn't supported.
- ISRs: Chapter 3 works, including those that use the SDA.
- Other: programs that work at all: (1) FNAM: FN\_2Fh AH=60h: REDOS PSP parent chain; (2) FNCSMD: INT 2Fh AH=Ah: installable command; (3) FNFDI: master environment; and (4) INTRPT: intercepting interrupts.

How can DOS be accessed by OS 2's set of set of valid DOS device drivers? If a DEVICE statement appears in CONFIG.SYS for a driver determined to be a DOS device driver, then the DEVICE or DEVICEHIGH is passed on, because DOS DEVICE settings, these settings apply to VDMs. If an MIVDM would load a DOS device driver in each VDM. Only character DOS drivers are supported. For example, to load ANSISYS into a VDM, you have a choice. Either place



DEVICE=C:\MS2\MDCS\ANSISYS in CONFIG.SYS so it will be loaded into all VDMs, or add C:\MS2\MDCS\ANSISYS to the DOS DEVICE setting for the particular VDMs to which you want it loaded.

OS 2 provides a VDDSetDeviceEx() function, which links a DOS device driver header into the chain of drivers to be passed to a DOS session at VDM creation. For example, the ANSISYS.VDD uses this function to force a fake VMSXXXX01.FDOS driver in the VDM address space. This is similar to the DOSMGR.AddDevice() function in Windows Embedded code, which we discussed in Chapter 1.

We asked one of the VDM developers at IBM to summarize what DOS programmers can do to help their programs run in DOS emulation environments such as OS 2:

In the 10 years I've worked in DOS Emulation I've seen an unbelievable number of bad programs. I've come up with a real quick, simple test of which is a DOS program vs. all behaved. On a DOS machine run the program. The load SHARE and try it again. You can't believe how many programs fail even on a generic DOS machine with SHARE loaded. The single biggest offender is INI 21h AX 31000+ open compatibility mode. Compatibility mode was a biggest mistake DOS 3.0+ all have possibly made. If I had to give DOS programmers suggestions on programming, number one on the list would be never use compatibility mode. Why? Compatibility mode changes meaning under different circumstances. A compatibility mode operation on EAX will not behave the same on a different EAX, or on no EAX.

Other suggestions:

Check return codes (i.e., the carry flag). You can't believe how many programs mess this up!

Avoid carrying a value that is mixed with other bits in the ESI/EAX. Next early DOS introduced some nasty bugs here.

Don't expect hard things to always work. For example, opening a file and, while it is open, trying to delete it, or opening a file with read/write access on a read-only media such as CD-ROM.

### New OS/2 Services for Old DOS Programs

Now that we have some idea of the support OS 2 provides for the standard DOS interface, what sort of new interface does OS 2 provide to DOS programs?

The OS 2 API, even as it exists, is not a first-class, portable, DOS program, being neither OS 2 nor a Windows API, yet is not as usable to DOS programs as a good Windows Low-level OS 2 does. Let's examine the extensions to INI 21h that make such a good portion of the OS 2 API accessible to DOS programs running in a VDM. First, these INI 21h extensions, a DOS program running under OS 2, could access some of the otherwise unavailable OS 2 kernel functionality described in IBM's OS 2 2.0C *Guide to Program and Programming Tools*:

First, there is the set of functions OS 2 has added as part of the Extended API (EAPI) mostly to support extended Attributes (EAs):

```
INT 21h AX=5702h DosQueryPathInfo() (1.1+) and DosQueryFileInfo() (1.2+)
INT 21h AX=5703h DosSetPathInfo() (1.1+) and DosSetFileInfo() (1.2+)
INT 21h AX=6C05h DosOpen2()
INT 21h AH=60h DosMkdir2()
INT 21h AH=60h DosEnumAttrib()
INT 21h AH=6Fh DosRmdirEASize()
```

The INI 21h AX=5702h and AX=5703h Extended Attribute functions are the most important. These allow a DOS program to indirectly call the OS 2 DosSetQueryPathInfo functions, which are described in detail in IBM's documentation. Interestingly, IBM's DOS 4.0 users may wonder these INI 21h Extended Attribute calls. To use the INI 21h calls, see the Introduction to the book.

any additional information can be inferred from Chapter 5 on Extended Attributes in IBM's *Control Point in Programming Guide*. As a side note, it's worth pointing out that any programs that hook the normal MS-DOS INT 21h AH=57h Cvt, Set File Date, and Time functions will need to explicitly ignore these OS/2 Extended Attribute functions.

JOIN and SUBST are provided to DOS boxes under OS/2, but they are implemented quite differently from their JOIN and SUBST in real mode DOS, where they muck around with DOS internals. Since JOIN and SUBST are built into the OS/2 file system as kernel services, the JOIN and SUBST commands do for OS/2 DOS boxes need to call into the kernel, using INT 21h AH=61h.

OS/2 INT 21h AH=61h JOIN/SUBST function

BH = 61h

AL = 0 (List), 1 (Add), 2 (Delete)

DX = drive number

CX = size of buffer

SI = 2 (JOIN), 3 (SUBST)

ES:DI → buffer

BP = 'dg' magic signature

DOS programs running under OS/2 can access named pipes using the normal DOS open, read, write, etc. calls using INT 21h AH=3Dh/40h calls on files such as PIPE\PipeName or \\SERVER\RPCPIPE\PipeName. For extended pipe support beyond the basics, there is a standard (though undocumented) LAN Manager-based network pipes interface:

INT 21h AX=5F32h	DosGetPipeInfo
INT 21h AX=5F33h	DosGetPipeHandState
INT 21h AX=5F34h	DosSetPipeHandState
INT 21h AX=5F35h	DosPeekPipe
INT 21h AX=5F36h	DosTransactPipe
INT 21h AX=5F37h	DosCallPipe
INT 21h AX=5F38h	DosWaitPipe
INT 21h AX=5F39h	DosRawReadPipe
INT 21h AX=5F3Ah	DosRawWritePipe

OS/2 2.00 itself uses these named pipe functions to implement DOS clipboard transfer and DOS dynamic data exchange (DDE). OS/2 3.1 instead uses the VWIN32S.DDD, which improves performance. For more information on these functions, see Mike Stuebs' article "The Undocumented LAN Manager and Named Pipe APIs for DOS and Windows" *Dr. Dobbs Journal* (Apr 1993).

OS/2 2.0 makes major extensions to DOS using INT 21h AH=64h. The handler for INT 21h AH=64h first looks for the C:\SYSTEM\DLL\64661.cpl, a developer's utility; then it's an OS/2 method. If not, then the handler jumps to the code for DOS extended open (INT 21h AH=6Ch). It's a good idea to look around for a bug in the API that used INT 21h AH=64h rather than AH=6Ch for the DosOpen2() function.

OS/2 INT 21h AH=64h handler finds the C:\ magic signature, it examines the BX register. If EBX is 0, it's a special request (see below). Otherwise, BX is the DPMSE API's ordinal number of an OS/2 API. DosXXX function to execute.

OS/2 API function calls Windows API functions are exported from dynamic link libraries (DLLs), and DDE export has an "ordinal number" given or decimal. For example, the OS/2 Dos32StartSession() API function's DPMSE API's 37. DPMSE API's is the module name for the OS/2 kernel, many of how DDE calls are exported provided by OS2KRN1, rather than by DPMSE API's DDE1, which is still the DPMSE API's ordinal number. INT 21h AH=64h expects a DPMSE API's ordinal number. It's not a valid ordinal number, which is why this indicates a special request. However, not every OS/2 DPMSE API's function is callable from a DOS program via INT 21h AH=64h; this interface supports only a specific set of OS/2 APIs.

INT 21h AH=64h BX=DOSCALLS ordinal (decimal) CX=636Ch

BX=37 (25h) Dos32StartSession (see below)

0x=130 (82h)	Dos32GetCP
0x=167 (A7h)	Dos32GetAttach
0x=191 (BFh)	Dos32EditName
0x=203 (CBh)	Dos32ForceDelete
0x=324 (144h)	Dos32CreateEventSem
0x=325 (145h)	Dos32OpenEventSem
0x=326 (146h)	Dos32CloseEventSem
0x=327 (147h)	Dos32ResetEventSem
0x=328 (148h)	Dos32PostEventSem
0x=329 (149h)	Dos32WaitEventSem
0x=330 (14Ah)	Dos32QueryEventSem
0x=331 (14Bh)	Dos32CreateMutexSem
0x=332 (14Ch)	Dos32OpenMutexSem
0x=333 (14Dh)	Dos32CloseMutexSem
0x=334 (14Eh)	Dos32RequestMutexSem
0x=335 (14Fh)	Dos32ReleaseMutexSem
0x=336 (150h)	Dos32QueryMutexSem (?)
0x=337 (151h)	Dos32CreateMutexWaitSem (*)
0x=338 (152h)	Dos32OpenMutexWaitSem
0x=339 (153h)	Dos32CloseMutexWaitSem
0x=340 (154h)	Dos32WaitMutexWaitSem
0x=341 (155h)	Dos32AddMutexWaitSem
0x=342 (156h)	Dos32DeleteMutexWaitSem (?)
0x=343 (157h)	Dos32QueryMutexWaitSem (?)

The most useful of these functions is `DosStartSession`, which a DOS program can use to start up other programs, including even OS/2 programs. The ability to launch OS/2 programs is not available with the `DOS EXECUTE` (INT 21h AH=4Bh) function, which will simply run the MZ portion of any file passed to it, so the following is quite useful:

```
DosStartSession
Entry
    AH = 64h
    BX = 37 (25h)
    CX = 636Ch ('cL' signature)
    DS:SI = STARTDATA structure
Exit
    AX = return code
```

The `STARTDATA` structure is the same one used by OS/2 programs that call `DosStartSession`, and includes the program name, the command-line arguments, environment, session type, an ESD file, screen coordinates, a position table, and so on. For the format of the `STARTDATA` structure, see Chapter 4, "Program Execution Control," of *HBM's OS/2 2.0 Cost of Program Programming Guide*.

In OS/2 2.1 LINK386 mode, INT 21h AH=64h `DosStartSession` call to implement a `exec` or `new` undocumented switch. `RU_NEROMVDM` Normally the OS/2 linker binds an OS/2 application with a DOS stub that just prints a "this program requires OS/2" error message and exits. `RU_NEROMVDM` does a stub that uses this INT 21h AH=64h `DosStartSession` call. Windows badly needs something like this. Microsoft's `WASRVREX` and `WVLE` present one solution to this problem, but really, this functionality needs to be built into Windows itself.

Returning to the other INT 21h AH=64h functions, if CX=636Ch, but BX is not a valid `DOSMATHS` ordinal constant, then the DX register indicates a special function. The Interrupt List on the accompanying disk describes these calls in more detail.

```
INT 21h AH=64h BX=0 CX=636Ch DX=subfunction
DX=0    Enable Automatic Title Bar Switch
DX=1    Set Session Title Bar
DX=2    Get Session Title Bar
DX=3    Get LASTDRIVE (used by WIN-OS/2)
DX=4    Get Size of PTDA JFF (used by WIN-OS/2)
```

```

DS=5  Get 05/2 SFT Second Flags Word (print spooling?)
DS=6  Unload DOSRNL Symbols and Load Program (debugging)
DS=7  Get I/O-05/2 Call Gate Address
DS=8  Get Loading Message (used by MSL)

```

As a final example of functionality that OS 2 provides to DOS programs, the OS 2 2.x VIDEO SWITCH CONTROL setting enables and disables the INT 2Fh AH=40h screen save API from OS 2 1.x. By default OS 2 doesn't issue the INT 2Fh AH=40h calls VIDEO SWITCH NUMBER VICON is off. Because issuing them has some implications for Super VGA and VGAhsaves, Windows also supports these screen switch functions, which are documented in the Windows 3.0k, see ENTERTAIN on disk for further details.

This implies, or at least hints, of the additional interrupt-based APIs that OS 2 provides to DOS programs. Remember, however, that any VxD can provide this sort of functionality to DOS programs, so using it in OS 2 VxDM is really what we've just looked at: what happens to come built in to OS 2. This is very similar to the situation with Windows Enhanced mode, where the additional APIs that VxDs can provide to DOS programs is potential, but not available even if it is certainly fairly limited. Someone just has to write a VxD in Windows or a VxD in OS 2 to make new APIs available. Of course, you also have to consider whether the benefits of the additional functionality are worth the additional overhead of the virtual device driver—remember DMSPPID in Chapter 3.

## DOS Emulation Under Windows NT

Windows NT is Microsoft's latest attempt to build a "real man's" operating system, with genuine 32-bit-based preemptive multitasking, demand-paged virtual memory, built-in networking, an improved file system, secure, group- and file-portability, and symmetric multiprocessing. Windows NT was not only on-line 80386 and higher processors but also on RISC processors such as the MIPS 4400. All these gross features require, at least eight, more realistically 12 or 16 megabytes of memory.

For its foreseeable future, NT will be a niche product. According to *Barron's* (April 5, 1993), Rick Pardo, a product manager for Windows NT at Microsoft, "says flat out that over the next two years 80% of the people using PCs will need the kind of computing power NT provides." Microsoft intends NT more as a replacement for MS-DOS or Windows, but as a way to perhaps conquer new markets, key clients owned by companies such as IBM and Novell. Even Microsoft's executive VP, Steve Ballmer, a knowledgeable man when NT is introduced, will at first be little more than a public relations man. *New York Times*, May 26, 1992. Mike Magnus, another Microsoft executive VP, more recently, tempering the company's not-so-great expectations for NT by noting, "If you don't know a company needs NT, then you don't need it." *Microsoft Soft Peddling Its Latest*, *New York Times*, May 24, 1993. In the next few years, perhaps fewer than 20% of existing Windows 3.1 users will move over to Windows NT.

As programmers considering whether to "bet the farm" on NT, this point is really important! Microsoft has been almost too explicit about how NT will play a relatively small role in the PC market place, or simply that how it will not replace Windows 3.1, and how it is intended especially for network server machines. Developers should therefore pay careful attention to what Microsoft is actually saying about NT, and not simply pick up on the amount they're talking about it. Frankly, Windows NT could turn out to be another OS 2 Yawn.

It is important to distinguish NT from the Win32 API, Microsoft's intended successor to the 16-bit Windows API. Whatever happens with NT, Win32 seems like a clear winner. Win32 is not at all tied to NT. A subset runs today on Windows 3.1 Enhanced mode, Win32s—and a larger subset will run on Chicago, DOS 7.0 and Windows 4.0—scheduled for release in 1994. In contrast to the prodigious memory requirements of NT, Microsoft's slogan for Chicago is "runs great in four megabytes."

Generally, when developers say they are writing an NT application, they really mean they are writing a Win32 application.

Just as Win32 applications don't run only on NT, conversely NT does not only run Win32 applications. Think of NT as a motherboard, into which you can plug multiple daughterboards. In NT terminology, these are subsystems. Win32 is just one NT subsystem; others are WOW, Win16, or Win32, responsible for running 16-bit Windows programs or NT, an OS/2 1.x subsystem, a POSIX subsystem, and a variant of MS-DOS subsystem called NTVDM.

As in OS/2 2.x, the same VDM means Virtual DOS Machine. Like OS/2 and Windows Enhanced mode, Natively compiled DOS programs execute in 8086 mode at the chip's lowest privilege level, Ring 3. As you can see, there are a number of similarities between the NTVDM subsystem and the VDM layer in OS/2 2.x. This is hardly surprising because, before Microsoft's divorce from IBM and its own formation of OS/2, Microsoft worked on OS/2 VMDM. The NTVDM team at Microsoft, led by Matthias Fritton, who, according to Helen Canter's *Inside Windows*, "had previously worked on OS/2's DOS compatibility circuitry in the Matt's house, roomy enough to provide much of the following material. Other material in his section comes from a talk by Sangeet Bhattacharya, "Virtual Device Driver Support on Windows NT," at Microsoft's NT device driver developer's conference, October 1992. Yes, like OS/2 2.x and Windows Enhanced mode, NT also has virtual device drivers, or VxDs. OS/2 2.x uses an array of VxDs rather than VxDs. However, unlike Windows VxDs and OS/2 VxDs, NT VxDs run in Ring 3 user mode, earlier than in Ring 0 kernel mode. NT has separate kernel mode device drivers that run in Ring 0.

### The Client/Server Model

As with VMDM in OS/2, the goal of NTVDM is to run DOS applications without DOS. Old DOS programs continue to have DOS INI 211 calls, but these are serviced using Win32 calls. In the remainder of this chapter, we try to make this "smoothed-out" path of applications explicit.

As an example, consider a simple program written in C that uses standard C runtime library functions such as `open`, `read`, `write`, and `close`. If this program is built with a C compiler for DOS, the resulting executable contains DOS calls such as `INT 21h`, `API 3Dh`, or `API 66h` to open the file `MyFile.txt` and `API 40h`, `system`, and `API 3Bh` to close. The corresponding DOS "MZ" executable could run on Win32 with a compiler for NT. This program, after recompiling for Win32 without changing a source code statement, still, and this is crucial, is a Win32 executable. The corresponding Win32 API calls for example, consist of using `INI 21h`, `API 3Dh`, or `API 66h` to open files, the recompiled Win32 version would contain calls to one Win32 CreateFile function. Yes, in Win32, existing files are opened with `CREATE_ALWAYS`. The ability to target multiple operating systems from a single source code is called *source compatibility*. It is not very difficult to achieve.

While extremely subtle, source compatibility is a nice compromise with which help several users. To satisfy this large and diverse segment of people, a new operating system must support many *compatibilities* with existing files for the popular operating systems. NT must be able to take old DOS executables containing INI 211 calls and low-level INI 211 calls as some sort of a signal that the application was built as a CreateFile, ReadFile, WriteFile, or Flush on Your local say if it a call to INI 21h 3Ah turns into a call to CreateFile.

There is not much source compatibility difference between INI 21h, API 3Dh, and CreateFile. Not only is DOS really a user-space DOS program running inside NT, it may not even be running on an Intel 80x86 machine. NT supports DOS emulation on RISC machines too. Even on Intel machines, it is hard disk management, not formatted using the old DOS file allocation table (FAT) format. NT also supports OS/2's large file format for systems, HPFS, and a new NT file system, NTFS. How can an old DOS program be made to work with all this new, non-DOS non-FAT possible new files, etc.?

Actually, this is not all that different from something that already takes place today under DOS. The DOS network redirector allows DOS INT 21h file I/O and directory calls to be serviced using files and directories located on other machines. As explained in Chapter 8, these other machines are probably not even running DOS. All a DOS program knows is that it issues the right INT 21h API 30 invocation to open a file and how and behold, it gets back a file handle. Most programs hardly care where the file handle came from. If the handle represents a file located on a NetWare file server (e.g., MIPN-P4000), a Sun Sparcstation, or a Macintosh, that makes no difference.

This network redirection sounds the same as what must go on when you run a DOS application on a non-DOS operating system, such as NT. And, except for the fact that the DOS client and the Win32 server may be located on a single machine, it is the same. NT is built around the client/server model. Rather than have the operating system provide APIs directly, all APIs are instead provided by services or servers, which have (or more or less) normal one-mode applications.

Supporting DOS INT 21h files or APIs with the idea, which NT borrows from the Mach operating system, that APIs and other facilities normally regarded as part of the operating system can instead be provided by user-mode applications. In other words, NT services a DOS application's INT 21h calls in exactly the same way that it services a Win32 application's Win32 calls with a subsystem. For a clear explanation of the NT client-server architecture, see Hector Custer's book *Inside Windows NT*, particularly the chapter on protected subsystems. For client-server operating systems in general and a breakdown of Mach, see Andrew Tanenbaum's brilliant *Modern Operating Systems*.

### NTVDM, NTIO, and NTDOS

To see a more detailed how NT runs DOS applications without DOS, possibly without the FAT file system, and possibly without an Intel processor, let's look at what happens when a user starts a DOS program under NT.

Windows NT runs programs with the Win32 CreateProcess API. When a user clicks on a program icon in the NT Program Manager, CreateProcess is called, even if the program is a DOS program or a 16-bit Windows application.

CreateProcess() is perhaps one of the lower-level undocumented functions that it invokes (see "Undocumented NT" and earlier chapters) inspects each FAT file to see the type of program it has been asked to run. All FAT files have two-byte signatures, magic, indicating their type: Win32 executables use Microsoft's Portable Executable format and have a PE signature. Both Win16 and OS/2 executables use the New Executable file format and have NE signatures; there is an additional image operation system flag in the executable header that distinguishes Windows and OS/2 programs. Both the NE and PE formats are supersets of the old DOS executable format, which starts with an MZ signature, the tribute of Microsoft's Mark Zukiowski. Thus, if CreateProcess() is handed a request for a DOS program, it MZ signature. CreateProcess() knows it is dealing with a DOS program. For some strange reason, for COM files, which have no header at all. An NE header with Windows as the target is a Win16 executable.

CreateProcess() deals with both DOS and Win16 executables by looking up a command line in the NT registry. By default, this command line tells CreateProcess() to instead run a Win32 executable, NTVDM EXE, and pass the name of the original DOS or Win16 program in which the user is actually interested. In other words, clicking on the DOS program DOSEXE results in the execution of the command "NTVDM EXE DOSEXE."

NTVDM EXE is not a DOS program but a user-level Win32 application. It creates a VDM in which a DOS program or Win16 program can run in VM86 mode. In this VDM, it loads and starts NTDOS.SYS in real mode, DOS.BK in address 0070:0080. NTDOS loads another file, NTDOS.SYS, and the process is CONHEX.SYS. If CONHEX.SYS contains HIMEM.SYS and a DOS-HIGH state, most work of the code is moved to the bottom of the VDM that lies just above one megabyte.

The names NTDOS.SYS and NTDOS.SYS should sound familiar. These are hacked versions of IO.SYS and DOS.SYS which form the core of MS-DOS 5.0 (see Chapter 6). While installed on NT, they are still real-mode DOS programs; these files represent the 16-bit side of the NT-DOS emulation layer. There is also a hacked copy of COMMAND.COM, which makes use of IOEXEC.BAT, the program that the user asked to run is loaded into this environment.

### Magic Pills and Bops

How do the  $N_{x}$  files NTDOS.SYS and NTDOS.SYS differ from the DOS files IO.SYS and MSDOS.SYS (or one of the MS-DOS files) as removed. What are the assumptions about the 386/486 system? As an example of the changes made to IO.SYS and MSDOS.SYS to create NTDOS.SYS and NTDOS.SYS, consider the following code fragment from MSDOS.SYS in DOS 5.0:

```
mov dx, cx:DOS_05
lds si, dword ptr ds:PTR_CURRCD5
test word ptr [si+DOS_ATTRIB], NEI_READIP      ; test [55943h], 8000h
```

The equivalent code in NTDOS.SYS looks like this:

```
mov dx, cx:DOS_05
lds si, dword ptr ds:PTR_CURRCD5
db 0C4h, 0C4h, 50h, 13h
```

On Intel processors, C4h C4h is an illegal instruction (it decodes as IFS AX,SP). Executing this illegal code in V86 mode causes a trap into whatever protected mode software is acting as the V86 monitor. In NT, this monitor is NTVDM.FAT.

Remember H.1 (DOS 7.2). A collection of the IBM engineers pointed earlier in this chapter described as OS/2's *magic pill* or *ARP* what is used for the same purpose. We do not need code? These magic pills instruct the real-mode software to enter V86 mode to protect itself. This is necessary because we use 16-bit real-mode code if we set our own access vector from 32-bit protected mode. We'll call C4h C4h the *N<sub>x</sub> DOS magic pill*.

C4h C4h is derived from the ISVOP. In the NT code driver for IOPL, BOP stands for BIOS operates, which is Microsoft's way of what people often call for this mechanism of how to 16-bit real-mode software. BOPs are call-backs to call down to 32-bit protected mode code such as NTVDM and NTVDDK.

ISVOP.H defines the following supporting third-party bops:

```
RegisterModule      db C4h, C4h, 58h, 0
UnRegisterModule    db C4h, C4h, 58h, 1
DispatchCall        db C4h, C4h, 58h, 2
```

Any DOS software running with NT can use these routines, which are documented in the NTVDDK *Win32 API Guide* (http://www.microsoft.com/windows/nt/ntvddk/). This can be done, too, even though with a 32-bit protected mode driver, it is about as NTVDDK as nothing more than DLLs and thus can call *Win32 API* to cross such as *CreateFile*, *ReadFile*, and so on. RegisterModule expects the ASCII name of the DLL and a Dispatch entry in the DLL. It returns a handle which can be used by DOS programs to then pass to Dispatch call to call to the 32-bit DLL. Clearly, this is a *Win32* program, not a *Win32* system.

Looking back at the code fragment for NTDOS.SYS, you can see that the string C4h C4h in the code is not calling any of the three factors from ISVOP. As noted earlier, ISVOP supports third-party bops. Microsoft uses its own bops that NTDOS.SYS, NTDOS.SYS, and COMMAND.COM use:

```
BOP_DOS      C4h, C4h, 50h      Used in NTIO.SYS, NTDOS.SYS
BOP_NOW      C4h, C4h, 51h      Win16 on Win32
BOP_XRS      C4h, C4h, 52h
```

BOP_DPMI	C4h, C4h, 53h	
BOP_CMD	C4h, C4h, 54h	COMMAND.COM to CMD.EXE?
BOP_DEBUGGER	C4h, C4h, 56h	
BOP_REDIRECT	C4h, C4h, 57h	Used to be 55h
BOP_REDPARITY	C4h, C4h, 58h	See ISVDDPH RegisterModule, etc
BOP_NOSUPPORT	C4h, C4h, 59h	Most warning dialog box
BOP_GRABTIMER	C4h, C4h, 5Bh	Check whether timer INTs required
BOP_VXD	C4h, C4h, 5Ch	temporary?
BOP_VXD2	C4h, C4h, 5Dh	temporary?
BOP_NOTIFICATION	C4h, C4h, 5Eh	16-bit to 32-bit notification
BOP_UNIMPINT	C4h, C4h, 5Fh	
BOP_SWITCHTOREALMODE	C4h, C4h, 60h	
BOP_UNSIMULATE	C4h, C4h, 6Fh	End execution of code in a VDM

I'm a little suspicious. Microsoft has thus defined its own little pseudo-like `bop` instruction set. NTDOS.SYS and NTVDM.SYS use these `bops`, especially BOP\_DPMI, essentially as VDM supervisor calls.

## Bops Through History

Microsoft's use of illegal instructions to make the hyperspace jump from 16-bit real mode to 32-bit protected mode sounds like an *ad hoc* hack that someone came up with in his own time. In fact, this `bop` technique of using an illegal instruction as a supervisor call dates back over twenty years to IBM's VM/360 operating system, where a nonex-ecutable instruction (often referred to as `DIAGNOSE`) was used to force an exception with in a virtual machine, so that the exception could be caught by the VM kernel and interpreted as a service request.

For more information on the history of this great hack, see the letters to the editor in *Windows Developer Journal* (July 1993) (pp. 99-100). These letters were prompted by a excellent earlier article of Neil VDDs, Paula Tomkinson, "Windows NT Virtual Device Drivers for Hardware Dependent 16-Bit Applications," *Windows DOS Developer's Journal*, May 1993, which included a short sidebar on bops ("The VDD Backdoor").

Windows 3.x and DOS software from third-party vendors uses bops to communicate with VDMs and even with NTDDK-like modules, see the chapter on "Virtual Device Drivers for MS-DOS Applications" that I've Special Illustrated. It shows how to use the RegisterModule and DeviceIoControl bops to connect a VDM that supports something like a LAN board or a 3270 communication adapter.

But bops are really, each more generally, those that that any DOS program could use bops to address the VDM, so maybe that's just Windows 3.11's or the DOS program's belief. If the bop is placed in a DOS file, it's not an ISR, the normal DOS programs that call the device driver or ISR or the bop just use bops to force a trap being aware of it. This transparency, of course, is the whole idea of bops, so they should be tracked as the 1225000 layer for a form of remote procedure call (RPC).

## What is NTVDM.EXE?

Microsoft's `bop` handlers are located in NTVDM.EXE, which, when we last left it, had started a VDM and loaded NTDOS.SYS and NTVDM.SYS. When these loaded versions of DOS's DPMI sub-systems are loaded, 50h (31) of the NTVDM.SYS code fragment above is handled in NTVDM.EXE.

NTVDM, like any PE executable, can be examined with a utility such as Microsoft's COFF EXE COFF stands for Common Object File Format, a UNIX standard upon which PE is based.



COFF EXE can even be run under DOS using Phar Lap Software's `LINE` 32-bit DOS extender that runs NE executables under MS-DOS. For example:

```
tot \ms-tools\bin\coff dump -reports -imports %ntvs386\ntvdm.exe
```

Actually, as this book was going to press, Microsoft replaced COFF EXE with LINK32 EXE. Examination of `NXVDM EXE`'s COFF shows that NXVDM exports a number of functions for use by VDDs, such as `VDDInstallHook`, `VDDInstallMemoryHook`, and so on. The `NXVDM` documents these functions. NXVDM also imports functions from `KERNEL32.DLL`, `USER32.DLL`, and `GDI32.DLL`. These dynamic-link libraries contain the functions documented in the Win32 SDK. Not surprisingly, Microsoft systems software after the NXVDM EXE also imports a number of real-protected functions from `KERNEL32`, such as `EXVDMGetNextVDMContext`, `RemoveContextVDM`, and `VDMContextOperation`, and uses an `NtVdmContext` handle to get out `NXVDM.DLL`. You will see later that `NXVDM.DLL` is *not* for the undocumented "native" NE API.

NXVDM EXE contains a number of components:

**DEMS2** is the slow-speed, protected-mode side of DOS emulation (DEM). When `NHOSYS` or `NTDOS.SYS` issue a trap, it comes to `DEMS2`.

**NXVDM** also contains a *native* (i.e., *execution mode*) `BEK`, which handles all the DOS program instructions such as `IN`, `OUT`, and `INT` that are illegal under DOS emulation. For example, `INT 21h` is an illegal instruction on Reg 3. `BEK` simply means that it is a instruction `0` because a trap into NXVDM which, in turn, can use the trap as an opportunity to do DOS emulation.

On non-Intel architectures, the DOS constraints on CPU emulation, which supports execution up to 80286 protected mode, and so on, around the Engineering Solutions Inc.'s `Soft9X` API. In recent different versions of NXVDM EXE, on different processors. For example, in addition to the other version of `386/NXVDM EXE`, there is also one for the R4000 MIPS NXVDM EXE. However, the exact same `NHOSYS` and `NXVDM.SYS` files used on Intel machines is used on MIPS machines.

As noted earlier, NXVDM exports functions for use as VDDs. These are documented in the `NXVDM Win32 User and Developer Guide` (see the chapter "Virtual Device Drivers for MS-DOS Applications That Use Special Hardware" in the *Book's User and Developer Release*).

NXVDM not only exports functions to use as VDDs, it also contains 100 VDDs for a stand-alone PC and 50 for real-time, remote video timer, timer, timer controller, and so on. New VDDs allow DOS applications running under NE to work, but this is going on I/O ports, writing to memory locations, and things of that sort that are normally activities. Meanwhile, they can also be executing on a PC or on other processor, although, a trap, means just that. The VDD using NXVDM functions such as `VDDInstallHook` and `VDDInstallMemoryHook` is to like the known low-level interface, the `ioctl` interface. The actual software, `ioctl` or `ioctl` can't be made up the `ioctl` of `ioctl` operations, the DOS applications can't be allowed to "pull" as Microsoft likes with its `ioctl` interface, even for files, even. An exception is a DOS application running for server, `ioctl` for `ioctl` will be allowed to direct bank emulation.

All these components, for OS 2 and Windows Enhanced mode. It is not an exception that these systems are limited, `ioctl` and `ioctl` of the machines. This is not an `ioctl` of a simulator, since the RISC machines or NE systems are probably have such a similar market share as to be a `ioctl` of the `ioctl`. In addition, NE VDDs are `ioctl` set mode, Reg 3 Win32 DLLs, not `ioctl` Reg 0 parts of the kernel, `ioctl` Win32 VDDs and OS 2 VDDs. Because they are `ioctl` mode Win32 DLLs, VDDs like NXVDM EXE can't use Win32 API, another difference from Windows VDDs which cannot make direct Windows API calls.

So how do all these pieces fit together?

NVDM.exe as a VDM starts NFDOS.SYS and NIDOS.SYS in the VDM and starts the DOS application that the user wants to run. The DOS application makes INT 21h calls, which go to NFDOS.SYS which supports only NVDM. NVDM's status VDI's, additional VDI's can handle any DOS-style disturbances that NVDM doesn't support. NVDM and VDI's are Win32 DLL's, so they can make Win32 API calls on the DOS program's behalf.

In some cases NVDM calls functions on KERNEL32.DLL. KERNEL32 contains not only the 32 bit standard functions of Windows API functions from KRNL386.EXE, but also a set of new Win32 API functions intended to replace the DOS functions that Win32 applications use. KERNEL32.DLL contains subroutines such as CreateFile, ReadFile, WriteFile, CloseFile, CopyFile, MoveFile, GetObjectInfo, SetCurrentDirectory, RemoveDirectory, DeviceIoControl, IOCTL, and so on. Note that these subroutines are the subroutines of the functions exported from KERNEL32.dll. For example, CreateFile uses CreateFileW, ASCII filenames, and CreateFileW, A wide - Unicode filename. NVDM uses KERNEL32 to carry out a DOS program's file I/O and directory operations.

KERNEL32.dll also contains a set of console functions to support character mode Win32 programs, such as Win95 console windows, like a console. NVDM uses the KERNEL32 Console to carry out a DOS program's input stream and stdout/stderr I/O.

Remember that the example of a DOS program that calls INT 21h API 310h or API-6Ch to open a file is not a DOS program. It is turned into a call to CreateFile. The "searching" part shown at the end of the example is a call to GetProcAddress, which helps determine NVDM.FAT, which then calls CreateFile on the DOS program's behalf.

But, since you probably called CreateFile, you may have a four byte file handle. This cannot be returned as a Win32 handle, which is expected to be a two byte number in AX. Instead, the four byte Win32 handle is stored in the NFDOS.SYS structure. Microsoft has modified from the DOS program, just as the DOS program itself is then turned to the DOS program's API associated with its %2. The FAT handle is a pointer to the DOS program's file handle. In other words,

```
MS-DOS: File handle -> JFF -> SFF
NTOS: File handle -> JFF -> modified SFF -> Win32 file handle -> ?
```

For NFDOS.SYS, the translation is more complex, since it has to be used in the OS 2.2's NVDM.EXE OS 2.10's "get handles" routine. On OS 3.11's SFI and MFI, a program, something similar happens. A NFDOS.SYS program might call NFDOS.SYS a DOS file I/O requests and let it is the necessary and use a file handle supplied from a Win32 application. A DOS structure such as the SFI structure is modified so that NFDOS.SYS is pointing down to NVDM, which in turn is calling CreateFile on KERNEL32 and so on.

## DOS 5.0

It is hard to tell how programs were written by other parts of this book, but are about the structure of SFI and other structures and DOS structures. Now, a lot of other take a few minutes to see what support NFDOS.SYS has for multi-tasking DOS programs.

A multi-tasking DOS program. The application is a console, so that it is not running as a native MS-DOS application. It is a console DOS application, but some will definitely find it this is not a DOS program. On the other hand, we saw that the 16 bit side of DOS emulates under NFDOS.SYS a multi-tasking DOS program, so the operation is quite close. Even many of the details of the program's structure are the same. NFDOS.SYS also has a lot of the same DOS based and DOS related functions and commands. NFDOS.SYS is also the FAT file system.

For the multi-tasking DOS side, any OS 2.2's NVDM advertises itself as DOS 2.0, so many NFDOS.SYS calls fail from the INT 21h API 30h and Version call. However, this caused many DOS programs to check for DOS 3.00 or 3.11 with error messages such as "This program can't run in a OS 2.0 DOS box". The same programs were happy to run when NFDOS returned a 5 x ver

sion number, so NTDOS now returns 511 from the INT 21h AH: 30h. This test sounds like a very bad move because it makes NTDOS difficult to distinguish from genuine MS-DOS SE. However, under NTDOS the function to get the genuine DOS version (INT 21h AX: 3300h) returns 50 (32h) as a patch against this. NTDOS currently (DOS 5.50) also has a version of the standard wired into the NTDOS version of COMMAND.COM.

What undocumented DOS functionality does this DOS 5.50 support? Microsoft has said that NTDOS will support only regular undocumented DOS functions but will not support DOS programs that depend on internal structures. It appears to be accurate.

INT 21h AH: 37h, of course, is supported, as is some of the SVS's data structure. For example, SWA's 21 program to the best of its ability checks for the VDM. The COMMAND program from Chapter 7 works very nicely under NTDOS, as does the WINDIV program from Chapter 3. Of course, these programs see only the memory chain for the first megabyte of the VDM.

Just as NTDOS provides an MR to boot, it has a boot DOS device chain with the NUI device rooted at the proper location in SWA's. As a result, the DIV program from Chapter 7 also runs nicely under NTDOS, as does a WINDIV program from Chapter 3. According to *Customs Inside Windows NT*, NT does some extra tricks to let DOS work device names on NT devices.

In the INT 21h AH: 30h example above, we noted that NTDOS modifies the SE1 to store four-byte Wans2 characters. In other words, NTDOS uses an SE1, but it is a SE1 of a different kind from the one in MS-DOS. In NTDOS, a file operation may succeed even if the file is on the Word2 side of the DOS separator line. The DOS SE1 structure only needs to contain 16 16-bit DOS state. As a result, some SE1 entries can be smaller in NTDOS. 71h bytes instead of 31h bytes. NTDOS does maintain the SE1 pointer in SWA's Running SE1 WALK, so using 3.5 indicates that NT automatically sets HIPS=255.

The CD is in NTDOS, as is standard. Different from the CD in MS-DOS, NTDOS ignores the LASTDRIVE value in SWA's and that in MS-DOS. ANSDRIVE holds the number of drives on the CD array. Eminent NTDOS maintains one CD for each physical drive and one for each of their mode drives, switching by single or dual CD is needed. As a result, the OKM program from Chapter 7, which tests that the ANSDRIVE value in SWA is the same as the ANSDRIVE value in DOS (INT 21h AH: 00h), reports success in NTDOS. Similarly, the FNT MDR program from Chapter 8 usually fails with a "no good CD" message because the value of ANSDRIVE is too low.

Without a proper CD, it is with a single value for ANSDRIVE, now all DOS network redirectors work. Certainly, the Phoenix program from Chapter 5 does not work under NT. (Perhaps Phoenix's state is dependent on the CD, that's most redirectors. NTDOS seems to use a VDM) DIR DIR and programs in directory 160h will fail.

How about ISRs? Here, it is easy to set up with all the ISRs from Chapter 9 work under NTDOS, including those for calls on INT 21h AH: 5100h and 51h swapable Data Arc. The only difference to an MS-DOS is NTDOS's global non-terminating handling of ISR masks. When an ISR terminates (except INT 21h AH: 31h or INT 21h), a message shows up regarding it. Windows NT means its resident program support. The message indicates that you can press Ctrl+ISR to work or press Ctrl+Esc to get a console and prompt. You are supposed to be able to hook a message on ISR from a DOS program, but this seems at first odd to work on an occasional basis.

In fact, ISR devices do work, not only in a console DOS application. When you are set up in a console, if prompt is N, it may look like a DOS prompt, but it is not. You are running CMD.EXE, which is a Win32 Console program. A VDM remains attached to this console window, but you are not using it now. When you see a DOS binary, the VDM waits for and does its thing. You can work only when a DOS application is actually running. CMD.EXE, which provides the C prompt, is not a DOS program. If you want continuous access to a TSR, then you should not press Ctrl+return to the C prompt, but get yourself a new Console window, running the command line dedicated for the TSR.

While COMMAND.COM is a Win32 program (i.e., an .exe) not only Win32 programs and, as you've seen, DOS programs, but also Win16 character-mode DOS 2 and DOS16 programs. Having one command meant that everything is a seamless departure from the lame DOS command prompt in Windows. However, not all DOS is made without a lot of extra tricks to run Windows programs. Microsoft calls COMMAND.COM "Service 1" and not "Service 2." At the same time, NT still needs the real-mode COMMAND.COM to process DOS batch files and generally to hang around in case DOS programs need it. For example, the undocumented COMMAND.COM hack for INT 2FH is available under NT/DOS, so even the HSI21 program from Chapter 10 runs very nicely.

Of course, the restricted DOS calls are just one example of odd behavior from DOS programs that is common to such as NT/DOS has to do to support, without compromising the integrity of the 4.0x sector. Another good example is the popular category 4 disk maintenance programs such as the Norton Disk Doctor EP. Tools. It's simply impossible for NT/DOS real-mode programs to write to the hard disk. You have a favorite DOS utility or debug utility you'll have to boot genuine MS-DOS to run it.

NT/DOS real-mode character DOS programs have direct *raw* access to the hard disk. The ratio is a 1:1 Win32/DOS application would expect to find FAT disk structures, whereas in fact there might be one NTFS or LFS. There may be a way in the future for NT to provide DOS programs with read access to the disks, but you'll see access to it of the question," says one of the NT/DOS developers.

Of course, the performance limitations of what even a native Win32 application can do as the user. Disk programs have the Same or Unimproved PC. Tools could be ported to Win32, but the user would have to be allowed to do as much to open the hard disk. Direct disk access is initiated with a call to `GetDiskInfo`, which only makes sense if you have a C: to drive C: 1 structure to write to the hard disk. To get a partition's volume ID, you'd use `GetDiskInfo` with `DeviceIoControl` (IOCTL\_DISK\_GET\_VOLUME). This is a low-level, raw operation. In fact, you can't do other operations on the drive. The paging swap file creates an open file handle, so you could not write or repair the disk with the swap file. Hammer cracks are welcome, but real NT fans may well demand MS-DOS for a while, if only to repair the massive hard disks NT needs.

### Additional NTDOS Functionality

It is possible that NT/DOS does not call for a lot of supporting existing DOS functions, including for reading and writing files, so you may see some of supporting DOS internal data structures, we need to work on some functionality for DOS programs. The DOS programs derive any positive benefits from running under NT. Certainly, DOS programs are going to run somewhat under NT than under genuine MS-DOS. But, in fact, there is a real reason for NT/DOS just a way to avoid having two machines on your desk or needing a DOS boot diskette!

We need to mention the COMMAND.COM (original command shell) and its shell to run a lot of other of applications. While it's not a true shell, what seems very important is that CMD allows different types of programs to run. COMMAND allows *pipe* between all the different types of programs it runs. This is a pipe that connects DOS and Win32 programs to a single command file. An interesting trick is to run the `Running DOS Examples` journal. John Richardson ("Escape from DOS to Win32"), Apr 1993, `etc`, says you can use Win32 pipes to use these anonymous pipes to request services of Win32 programs.

Of course, a Win32 application, which, as we've seen in Chapter 3, is really just a fancy form of process and mode, DOS application and DPMI client, can interact with Win32 applications using the `Clipboard` and `Clipboard` data exchange. Unfortunately, NT does not allow any further mixing. For example, Win32 programs cannot directly use Win16 DLLs, and Win16 programs can not directly call Win32 DLLs. The word "free" is important here because numerous indirect methods, such as named shared memory and DDE, are available.

While the NT services to DOS programs are rather sparse, it's important to remember the third-party traps that Microsoft supports. Using the RegisterModule and DispatchIO trap mechanisms, real-mode DOS code running under NT can call down to a VDD, which in turn can call any arbitrary Win32 API function on the DOS program's behalf. The NT DDK CD-ROM provides some good examples in the APPINITS.COM VDD, DOSMIO.CHI, and ADMINITS sample programs.

On the other hand, if someone is prepared to modify their DOS program to use bopps and is prepared to write a VDD to act as a Win32 surrogate for the DOS program, they should probably bite the bullet and port their program to Win32. The third-party traps are probably not as useful in a limited number of cases where the DOS software simply cannot be ported to Win32, but where there is an opportunity to either modify a small portion of it or to supply a DOS device driver or TSR that bops down to a VDD.

### Undocumented NT

The goal of this chapter has been to discuss "other DOSs," avoiding major discussions of any other aspect of NetWare OS 2 or NT. But in the case of Windows NT, it's difficult to resist a small peek at the broader topic of undocumented NT functions. Not only because your readers and Microsoft too have asked us why (*Undocumented NT* is selling out), but because the topic of undocumented NT helps clarify the nature of the NT operating system.

In contrast to the NT, Microsoft has provided massive amounts of programmer's documentation. This documentation starts quite early that its subject is the Win32 API, that is, functions such as (to pick a few at random) `CreateLocalMemory`, `MoveWindow`, and `LoadLibrary` exported from `KERNEL32.DLL`, `USER32.DLL`, `GDI32.DLL`, and other Win32 dynamic link libraries.

They're not the NT API. To repeat the point made earlier, Win32 and NT are not the same thing. Large subsets of the Win32 API run on platforms other than NT, for example, on Windows 3.1 Enhanced mode, Windows 95, and 4 Images OS. NT Win32's real purpose is a system. And NT supports APIs other than Win32, too. DOS, for example, we looked at in a previous chapter. Win32's a bit privileged in that, as we saw, it's not a superset of other subsystems. It is also a superset, in that the NT integration was a view of the other subsystems as optional and only back if you load an application that needs them. But in many ways, Win32 really is a NT in much the same way that DOS emulation does. It's a subsystem.

The beauty of NetWare is that any features that can be associated with the operating system kernel, such as API provision, have been moved into more or less normal applications, which can be preempted, pageloaded on other machines, and so on. It's not at all clear how many people really need this, but it's certainly cool.

So if Win32 and the other subsystems are all user-level code, where is the operating system kernel? It'd be nice to see a `KERNEL32.DLL` which really is combining enough system kernel mode but user mode code. We saw that the `NTVDM` subsystem is implemented in part using `KERNEL32.DLL`, but how's `KERNEL32.DLL` itself implemented. And what was that `NTDLL.DLL` file that `NTVDM.EXE` also used?

Much of the general-purpose NT API is exported from `NTDLL.DLL`, for example, Helen Carter's excellent *Inside the Zone: NT* mentions many NT objects, such as Processes, already exist, Events, and so on, and refers to the types of operations new objects provide, such as `Create`, `Open`, `Close`, `Query`, and `Wait`. But, understandably, the book doesn't discuss where to find these objects and operations.

If you examine `NTDLL.DLL`'s export dump, for example, you will see that `NTDLL` exports many of the NT functions whose existence one would infer from *Inside Windows NT*. For example, here's a small portion of the `NTDLL` export table:

61	62	0000b810	NtConnectPort
62	63	0000b820	NtContinue
63	64	0000b830	NtCreateDirectoryObject
64	65	0000b840	NtCreateEvent
65	66	0000b850	NtCreateEventPair
66	67	0000b860	NtCreateFile
67	68	0000b870	NtCreateKey
68	69	0000b880	NtCreateMailslotFile
69	70	0000b890	NtCreateMutant
70	71	0000b8a0	NtCreateNamedPipeFile
71	72	0000b8b0	NtCreatePagingFile
72	73	0000b8c0	NtCreatePort
73	74	0000b8d0	NtCreateProcess
74	75	0000b8e0	NtCreateProfile
75	76	0000b8f0	NtCreateSection
76	77	0000b900	NtCreateSemaphore
77	78	0000b910	NtCreateSymbolicLinkObject
78	79	0000b920	NtCreateThread
79	80	0000b930	NtCreateTimer
80	81	0000b940	NtCreateToken
81	82	0000c040	NtCurrentTeb
82	83	0000b950	NtDelayExecution
83	84	0000b960	NtDeleteKey
84	85	0000b970	NtDeleteValueKey
85	86	0000b980	NtDeviceIoControlFile
86	87	0000b990	NtDisplayString
87	88	0000b9a0	NtDuplicateObject
88	89	0000b9b0	NtDuplicateToken

You won't find any of these functions in the Win32 SDK or even in the NT SDK documentation—some are mentioned as deprecated. DOI, INC, NIDDI, and NSTATUS. They really are undocumented. On the other hand, most of these functions do have documented equivalents in the Win32 API, so it's not clear how important undocumented NT really is to most programmers. For example, in early beta versions of NT, Microsoft's PVIEW (Process Viewer) utility relied heavily on functions in NIDDI, DDI, and some March 1993 beta, much of the functionality upon which PVIEW relies had been incorporated into the Win32 API.

So, even if every NIDDI, DDI function has equivalent functionality in Win32—which is not the case—it's probably no use at this time because it helps send by much of the material Custer presents in her book.

So is NIDDI, DDI the genuine core of NT? Hardly. Looking back at the excerpt of the NIDDI export table, you may notice something odd: The third column shows the virtual address for each function, most of the functions are only 108 bytes apart. Microsoft sometimes likes to brag about its "high code" but it's unlikely that was implemented each NT function in only 16 bytes. Disassembly of a most important portion of NIDDI shows what's really going on. For example:

```
NtCreateProcess proc far
    mov eax, 1Ah
    lea edx, dword ptr [esp+4]
    int 2Eh
    retn 20h
NtCreateProcess endp
NtCreateProfile proc far
    mov eax, 1Bh
    lea edx, dword ptr [esp+4]
    int 2Eh
    retn 1Ch
NtCreateProfile endp
NtCreateSection proc far
    mov eax, 1Ch
    lea edx, dword ptr [esp+4]
```

```
int 2Eh
retm 1Ch
NtCreateSection endp
```

It goes on like this for hundreds and hundreds of functions. Most of the functions in NTDLL.DLL are little more than wrappers around INTEL instructions with function numbers in EAX.

Just as there needs to be a way for 16-bit real mode x86 running in a DOS application to make a SVC-like call down to the 32-bit protected mode subsystem code in NTVDM, there also needs to be a way for code running in user level (Ring 3, 32-bit protected mode code) to make a SVC-like call down to the NT kernel running at Ring 0. NTDLL uses INTEL instructions to call down to the NT kernel. Naturally, this protected mode INTEL instruction has nothing to do with the INTEL instruction real mode COMMAND.COM provides as its backdoor.

When NTDLL uses an INTEL instruction, it is handled by NIOSKRNL.EXE, the NT operating system kernel, or by NTKRNLMP.EXE, the multiprocessor version of the NT kernel. These files too can be examined with COFF-DUMP or a similar utility. NIOSKRNL exports almost 600 functions, such as (to pick a bunch at random):

```
CcGetFileObjectFromCcb()
ExQueueWorkItem()           ; Ex = General executive routine
FsRtlRegisterUncProvider()  ; Fs = file system
IoAdapterObjectType()      ; Io = I/O subsystem
IoStartNextPacket()
KeWaitForMultipleObjects() ; Ke = Kernel
KeMapIoSpace()             ; Ke = Memory management
NtCreateFile()
NtVdmControl()
ObQueryNameString()        ; Ob = Object management
PsThreadType()             ; Ps = Process structure
RtlGetHeapUserValue()      ; Rtl = Run-time library (user mode)
SeOpenObjectAuditAlarm()   ; Se = Security
ZwReadFile()                ; Zw = Windows NT system service
```

That is what the NT API looks like. It is important to note that the NT DLL *kernel mode Driver Reference and Kernel-mode Driver Design Tools* document some of these NIOSKRNL functions, and that some also are exported in DLL header files, such as NTDLL.H.

NIOSKRNL is part of the system HAL DLL, which contains the NT Hardware Abstraction Layer (HAL) interface port to the system NIOSKRNL. Despite the picture painted in most architectural diagrams of Windows NT, HAL is far from being the only hardware-dependent piece of NT; this is simply a goal.

NT internals is a major topic, and if NT goes anywhere, it deserves its own book, so we will leave the reader dangling here and return to the topic of undocumented MS-DOS.







IBM as part of its Personal Development Software series, is still available; contact Personal Development Software, Wallingford CT.)

What sets INTRNPI apart from other DOS monitoring programs is its use of a *scripting language* for intercepting interrupts. INTRNPI knows very little about any particular DOS or BIOS call. It doesn't know that INT 21h Function 3Dh is the Open File function, it doesn't know that this function takes the ASCII path name of a file to open to in the DS:DX register pair or that (successful) the function returns a file handle in the AX register. Rather than hard-wire such protocol knowledge into the program, INTRNPI's scripting language lets the user provide such knowledge.

The benefit is that the program is open-ended. If you want to monitor some undocumented region of DOS that this book doesn't mention, you can. If you want to examine some little-known DOS file system, you can, just write a script. Furthermore, because the INTRNPI language includes support for logical structures, you can produce meaningful output rather than raw register dumps.

INTRNPI is available for real-mode DOS only. To watch interrupts occurring in protected mode under Windows, you can use the BPINT command in SoftICE Windows or INTRNPI's baby brother, *Win-WINPI* (Windows I Spies), in Chapter 4 of *Undiscovered Windows*.

## A Guided Tour

Let's begin with a quick session with INTRNPI. It makes sense to start with something that does a little, so understanding DOS, so we will pretend we are interested in tracking down who tries to register queries. One could, of course, disassemble the program or run it under a debugger, but it makes more sense to treat the program as a "black box" and sample see which DOS file calls it makes. In other words, in this case we should study the program as a behavioral rather than a structural, job. In short, INTRNPI is perfect for this sort of exploration.

For example, the following command could be used to find out which files Microsoft's MSBROWSE is compiling, the non-industry standard HTML04 program:

```
intrnpi
cmdspy compile fopen scr
ci http://c
cmdspy report
```

This command loads INTRNPI, which is a memory-resident program. It then uses CMDSPY to intercept INTRNPI's scripted FOPEN.SCR. As explained below, CMDSPY communicates with the resident INTRNPI program through the Microsoft C! program's run. Finally, the cmdspy produces a report, which we also could have seen via a DOS file redirection, as in CMDSPY REPORT > C:\LOG. As for the FOPEN.SCR script itself, Listing 5-1 shows a very simple version which traps only calls to INT 21h Function 3Dh (Open File).

### Listing 5-1: FOPEN.SCR (simple version)

```
* FOPEN.SCR
intercept 21h
function 3dh      , Open File
  on_entry
    output "OPEN " (ds:dx->byte,asciz,64)
  on_exit if (cf == 1)
    samefile " [FAIL - ax =]"
```

This script instructs INTRNPI to intercept INT 21h and trap all calls to AH=3Dh (Function 3Dh, the DOS handle-based file open function). On entry to the call, INTRNPI should output the string "OPEN " plus the ASCII string pointed to by the DS:DX register pair; a maximum of 64 bytes will be stored. The pointer operation `(ds:dx->byte,asciz,64)` indicates that DS:DX is to be used as a pointer to an area of memory that should be treated as bytes. Sixty-four bytes should be stored and formatted on output as

an ASCII string. On exit from INT 21h AH 30h, if the carry flag is set, the string is INTRSPY it should output, on the same line, the string " FAIL " and the value of the AX register enclosed in square brackets. Note in Figure 5-1 how the on\_entry clause corresponds to the parameter expected by INT 21h function 30h, and how the on\_exit clause corresponds to its possible error return value.

**Figure 5-1: The DOS Open Function and a Corresponding INTRSPY Script**

<pre>INT 21h Function 30h Open File Call with:     AH = 30h     AL = access mode     DS:DX -&gt; seg.offs of ASCII pathname Returns:     CARRY = clear if function successful            AX = handle     CARRY = set if function unsuccessful            AX = error code</pre>	<pre>intercept 21h function 30h on_entry output "OPEN "            (ds dx-&gt;byte,ascii,64) on_exit            if (cflag = 1)                sameLine " [FAIL " ax "]"</pre>
--	---

The output from INTRSPY goes not directly to your screen, but into a results buffer. The default buffer size is 5K. This can be overridden by the INTRSPY command line, for example INTRSPY /10000. After running CLHELLO.C, the command CMDSPY REPORT produced the results shown in Figure 5-2.

**Figure 5-2: INTRSPY FOPEN Results for CLHELLO.C**

```
OPEN \msc700\bin\CL_EXE
OPEN \msc700\bin\ms32krm.dll
OPEN CON
OPEN CON
OPEN C:\msc700\bin\CL_EXE
OPEN C:\UNDOC2\CHAPS\13216_EXE [FAIL 2]
OPEN C:\msc700\bin\CL3216_EXE
OPEN C:\UNDOC2\CHAPS\015592sy [FAIL 2]
OPEN C:\UNDOC2\CHAPS\015592sy
OPEN C:\UNDOC2\CHAPS\015592st [FAIL 2]
OPEN C:\UNDOC2\CHAPS\015592st
OPEN C:\UNDOC2\CHAPS\015592ex [FAIL 2]
OPEN C:\UNDOC2\CHAPS\015592ex
OPEN C:\UNDOC2\CHAPS\015592in [FAIL 2]
OPEN C:\UNDOC2\CHAPS\015592in
OPEN C:\UNDOC2\CHAPS\015592st [FAIL 2]
OPEN C:\UNDOC2\CHAPS\015592st
OPEN C:\UNDOC2\CHAPS\hello.c
OPEN C:\UNDOC2\CHAPS\stdio.h [FAIL 2]
OPEN C:\msc700\include\stdio.h
OPEN C:\UNDOC2\CHAPS\q23.exe [FAIL 2]
OPEN C:\msc700\bin\q23.exe
OPEN C:\msc700\bin\q23.exe
OPEN 015592g
OPEN 015592ex
OPEN 015592in
OPEN 015592st
OPEN hello.obj [FAIL 2]
OPEN hello.obj
OPEN 015592sy
OPEN C:\UNDOC2\CHAPS\015592ln [FAIL 2]
OPEN C:\UNDOC2\CHAPS\015592ln
OPEN C:\UNDOC2\CHAPS\link.exe [FAIL 2]
OPEN C:\msc700\bin\link.exe
OPEN C:\msc700\bin\link.exe
OPEN C:\msc700\bin\link.exe
```

```

OPEN C:\msc700\bin\link.exe
OPEN C:\msc700\bin\link.exe
OPEN Q155921x
OPEN hello.obj
OPEN \msc700\lib\OLDNAMES.LIB
OPEN \msc700\lib\SLIBCE.lib
OPEN \msc700\lib\OLDNAMES.LIB
OPEN C:\hello.exe [FAIL 2]
OPEN C:\hello.exe
OPEN C:\UNDOC2\CHAPS\CROSPY.EXE

```

With a few extra simple HEURISTIC MIS2KRN1.DIF, a hacked Win32 portable executable PE, and a MSC00.SYS. The temporary files with suffixes like msc7p, and bk suffixes relate to different compilation stages. It was especially global optimization, local optimization, and so on.

It's a pity that Microsoft could not try the /q option to do a quick compile. Instead of the /q option, we could have used a full optimizing compiler, the INTRSPY results, as shown in Figure 5-3, are the result. Now under the /q switch, linker unfortunately, it does not have the same effect as MSC00.

**Figure 5-3 INTRSPY FOPEN Results for CL\_QC\_HELLO.C**

```

OPEN gcc.exe [FAIL 0002]
OPEN hello.obj
OPEN hello.c
OPEN c:\c600\include\stdio.h
OPEN 0072811k [FAIL 0002]
OPEN link.exe [FAIL 0002]
OPEN c:\c600\bin\link.exe
OPEN 0072811k
OPEN hello.obj
OPEN c:\c600\lib\SLIBCE.lib
OPEN C:\hello.exe
OPEN C:\hello.exe
OPEN C:\DOS\COMMAND.COM

```

With a few extra steps, this Win32 version INTRSPY script, we have created a file open spy, that is, a file watch. It now may take many more lines of code, and more important, a good several hours of programmer's time to write a linking C++ assembly language or Pascal.

Now, let's do a NLSR file system implementation. For a logging utility, for example, a file system (FSD) or file system, we are going to start lots of file system activity as DOS first tries to find and create a file. For a better idea, the log files I show are files being created. Now does it look for a file to be created? It does open files, such as IN, 2, 1, ALL, OBJ, FC, B, Open, and ALL, 0, 0, h. Extended Open, Create.

As you can see, it is a trap some additional functions, as shown in the boxed up version of FOPEN.NLSR. Listing 5-2 will show the DOS command lines with which programs are EXEC'd. Of course, since INTRSPY.NLSR is a FILE, that represents the parameter block used by the DOS EXEC function, it uses a ENTRY.CURL to represent the part of the DOS File Control Block that it's interested in.

**Listing 5-2: FOPEN.SCR Watches File Open, Create, Find, and Exec**

```

; FOPEN.SCR
structure param_blk fields
    env_seg (word,hex)
    args (dword,ptr)
structure fcb fields
    drive_num (byte,char)
    filename (byte,char,8)
    ext (byte,char,8)

```

```

; etc.
intercept 21h
function 0Fh ; Open File with FCB
  on_entry
    output FCBOP " (ds:dx->fcb.filename) ". (ds dx->fcb wxt)
  on_exit if (cf == 0Ffh)
    sameline " [FAIL]"
; -----
function 3ch ; Create File
  on_entry
    output CREAT (ds:dx->byte,asciiz,64)
  on_exit if (cf == 0Ffh)
    sameline " [FAIL (ax, dec) ]"
; -----
function 3dh ; Open File
  on_entry
    output "OPEN " (ds:dx->byte,asciiz,64)
  on_exit if (cf == 1)
    sameline " [FAIL " (ax, dec) "]"
; -----
function 4bh ; Execute Program
  subfunction 00h
  on_entry
    output "EXEC "
      (ds:dx->byte,asciiz,64) ; program
      (es:bx->param_blk.args->byte,pascal,128) ; cmdline
  on_exit if (cf == 1)
    sameline " [FAIL " (ax, dec) "]"
; -----
function 4eh ; Find First File
  on_entry
    output "FIND " (ds:dx->byte,asciiz,64)
  on_exit if (cf == 1)
    sameline " [FAIL " (ax, dec) "]"
; -----
function 6Ch ; Extended Open/Create
  on_entry
    output "XOPEN " (ds:dx->byte,asciiz,64)
  on_exit if (cf == 1)
    sameline " [FAIL " (ax, dec) "]"

```

The first INTRSPY STRUCTURE statement corresponds to the first two fields (ds:dx) only; the matter here is of the parameter block that INT21h Function 4Bh Subfunction 00h expects from the FCB register pair. FCBOP is output, the DOS command line of the first 64 bytes of it, with this expression:

```
(es:bx->param_blk.args->byte,pascal,64)
```

This indicates that the FCB pair points to an area of memory that should be treated as a param blk structure, and that we are situated in the argv field. FCBOP then uses another arrow (>) to indicate that args, which was defined as dword hex, points to a string. A Pascal designation, which reflects how strings are stored in Pascal, is different from ASCII, in that its first byte is a length count. This corresponds exactly with the command line used by MSDOS.

Now, of course, INTRSPY produces even more output, a small fragment of which is shown in Figure 5-4, describing the activity generated by the simple command CLHELLO.C.

**Figure 5-4: Detailed INTRSPY FOPEN Results for CLHELLO.C**

```

FIND cl.??? [FAIL 18]
FIND \bin\cl.??? [FAIL 18]
FIND \dos\cl.??? [FAIL 18]
FIND \bor.and\bin\cl.??? [FAIL 18]

```

```

FIND \watcom\bin\cl.??? (FAIL 10)
FIND \msc700\bin\cl.???
OPEN \msc700\bin\CL.EXE
EXEC \msc700\bin\CL.EXE hello.c
OPEN \msc700\bin\ms32brn1.dll
OPEN CON
OPEN CON
    Lots of OPENS, as in Figure 5-2
OPEN C:\msc700\bin\iq23.exe
OPEN C:\msc700\bin\iq23.exe
EXEC C:\msc700\bin\iq23.exe
OPEN 015592gl
OPEN 015592ex
    more OPENS, as in Figure 5-2 .
OPEN C:\msc700\bin\link.exe
OPEN C:\msc700\bin\link.exe
EXEC C:\msc700\bin\link.exe @015592ll"
OPEN C:\msc700\bin\link.exe
OPEN C:\msc700\bin\link.exe
    yet more OPENS, as in Figure 5-2
OPEN \msc700\lib\OLDNAMES.lib
OPEN C hello.exe
FIND caddsp.???
CREAT hello2.log
OPEN C:\UND0C2\CHAP5\CWDSPY.EXE
EXEC C:\UND0C2\CHAP5\CWDSPY.EXE report

```

The first thing Figure 5-4 shows is that before the actual execution of `CL.FIX`, `absP`, `COM`, `MAND`, `COM` must first find it. The series of failed calls to the `DFS` and `Find` function show `COM`, `MAND`, `COM` looking through its `PATHT` in many subdirectories before it finally finds `CL.FIX`. If you were going to be using Microsoft C, for it would probably be a good idea to optimize your `PATHT` by moving `C:\MSVCBIN` forward a little.

`INTRSPY` can be likened to a protocol analyzer for PC software interrupts, where `INTRSPY` itself only knows about raw interrupts, registers, and encryptions, and where the user's scripts impose the necessary higher level interpretation to understand what is actually going on.

### Device Drivers

`INTRSPY` also provides a script-driven approach to watching block and character mode device driver calls using the same `INTERCEPT` command used to hook interrupts.

```

INTERCEPT STRATEGY(device) ; Character mode device
INTERCEPT INTERRUPT(device) ; Block mode device

```

Block mode `IO`s always support data transfers to and from devices a block at a time, usually for random access storage devices such as hard disks, and are identified by a drive letter. Character mode device drivers provide stream character access to devices such as communications channels. They are also used to provide interfaces to system facilities such as XMS that may need to get notified early on during system initialization. DOS includes four in-block mode device drivers for the floppy and hard disk drives, for example, while `THEMFSYS` implements a character mode device driver for monitoring XMS extended memory. The names of these character mode drivers are displayed by programs such as `DEV` from Chapter 7.

Figure 5-3, `DDSR`, shows how you might record device driver calls to any device driver. Note that while you can monitor calls to both the `Strategy` and `Interrupt` entry points to the device driver, generally it is useful to monitor the `Strategy` routine. DOS does not exploit the separation in drivers between `Strategy` and `Interrupt` as soon as DOS calls the `Strategy` routine, it calls the `Interrupt` routine. Unlike operating systems such as OS/2 and NE-DOS, DOS does not provide overlapped, asynchronous I/O.

**Listing 5-3: DD.SCR Watches Calls to Device Drivers**

```

; DD.SCR -- watch any device driver named on command line
; CMDSPY COMPILE DD C
; CMDSPY COMPILE DD EMRXXXX

structure reqv fields
    len (byte, hex)
    subunit (byte, hex)
    cmd (byte, hex)
    status (word, hex)
; other stuff depends on packet

intercept strategy( "I" )
on_entry
    if ((es:bx->reqv.cmd) == 0) output "00 - init"
    if ((es:bx->reqv.cmd) == 1) output "01 - media check"
    if ((es:bx->reqv.cmd) == 2) output "02 - build bpb"
    if ((es:bx->reqv.cmd) == 3) output "03 - ioctl input"
    if ((es:bx->reqv.cmd) == 4) output "04 - input"
    if ((es:bx->reqv.cmd) == 5) && (ES == 0) ; often a lot of these
        output "05 - nondestruct input"
    if ((es:bx->reqv.cmd) == 5)
        incr ES
    if ((es:bx->reqv.cmd) == 6) output "06 - input status"
    if ((es:bx->reqv.cmd) == 7) output "07 - input flush"
    if ((es:bx->reqv.cmd) == 8) output "08 - output"
    if ((es:bx->reqv.cmd) == 9) output "09 - output w/verify"
    if ((es:bx->reqv.cmd) == 0Ah) output "0A - output status"
    if ((es:bx->reqv.cmd) == 0Bh) output "0B - output flush"
    if ((es:bx->reqv.cmd) == 0Ch) output "0C - ioctl output"
    if ((es:bx->reqv.cmd) == 0Dh) output "0D - device open"
    if ((es:bx->reqv.cmd) == 0Eh) output "0E - device close"
    if ((es:bx->reqv.cmd) == 0Fh) output "0F - removable media"
    if ((es:bx->reqv.cmd) == 10h) output "10 - output until busy"
    ; 11h and 12h are missing from DOS 5 Programmer's Reference
    if ((es:bx->reqv.cmd) == 11h) output "11 - stop output"
    if ((es:bx->reqv.cmd) == 12h) output "12 - restart output"
    if ((es:bx->reqv.cmd) == 13h) output "13 - generic ioctl"
    if ((es:bx->reqv.cmd) == 14h) output "14 - unused"
    if ((es:bx->reqv.cmd) == 15h) output "15 - unused"
    if ((es:bx->reqv.cmd) == 16h) output "16 - unused"
    if ((es:bx->reqv.cmd) == 17h) output "17 - get logical device"
    if ((es:bx->reqv.cmd) == 18h) output "18 - set logical device"
    if ((es:bx->reqv.cmd) == 19h) output "19 - act: query"
    if ((es:bx->reqv.cmd) > 19h) output ((es:bx->reqv.cmd) - UNKNOWN)
; could also handle IO-ROR driver commands

; Intercept interrupt("I")
; on_entry
;     output "Interrupt entry point polled"

run "X2 X3 X4 X5 X6 X7 X8 X9"
report

```

This script introduces some additional features of the INTRSPY language. The "I" is the script operator, and the word "I" would be a back tick if it's a parameter specified on the command line and is not a part of the script when it is compiled. The RUN statement is equivalent to running a command on the DOS command box, except that in the INTRSPY script, RUN is a flag that reduces the amount of time that you spend in programs you are currently interested in. CMDSPY runs the command lines on the fly, so you're compiling the script. In all, the REPORT statement is similar to running CMDSPY REPORT.

To see what commands are sent to the A: disk driver during a simple command such as COPY A:, you could issue the following at the DOS command prompt:

```
C:\ND00C2\CHAPS>cmdspy cnpfile dd a "command /c vol a:" > tap.tap
```

This keyboard-stroke command compiles DD.SCR to watch drive A. Issues the internal command, VOL A, issues the COMMAND.COM, and redirects CMDSPY's output to TAP.TAP. The output would look something like that shown in Figure 5-5. The first line is output from the VOL command itself; the rest of the output comes from CMDSPY.

**Figure 5-5: DD.SCR Output for VOL A:**

```
Volume in drive A has no label
(Handler for calls to 0070 06F5 was already stopped )
File, DD.SCR
Line 12: intercept strategy("a")
***
Warning: Same block device driver supports multiple drives
1214 bytes of code generated for "a" driver strategy script

Running program: C:\DOS\command.COM. .
Run completed.

D5 - nondestructive input
D1 - media check
D4 - input
D4 - input
D4 - input
D4 - input
D4 - input
D4 - input
D4 - input
D4 - input
D4 - input
D4 - input
D1 - media check
D4 - input
13 - generic ioctl
D1 - media check
D4 - input
D4 - input
D4 - input
D4 - input
D8 - output
D8 - output

Counter #5 : 2
All other counters zero.
```

Another example of INTRNRY that DD.SCR uses is a counter. There will often be so many calls to a device that it's not possible to do nondestructive input; that the INTRNRY results buffer would immediately fill up with thousands of nondestructive lines of output for each one. Therefore, DD.SCR only displays the first nondestructive computer responses at the rest only to running of counter INCR #5 in Listing 5-3.

Various counters are also identified by #1 through #16. Each is a 32-bit value that may be incremented, INCR, decremented, DECR, or reset to zero, ZERO, within a script. They may also be compared to test statements and used as elements in CPU, CPU, or SAMETIME statements. We pick #5 here simply because it matches the counter number for nondestructive input. At the end of an INTRNRY report, CMDSPY displays active counters. In Figure 5-5 we can see that there were only two active devices; in some cases, there will be thousands.

With many character-mode device drivers such as AM5XXXX0 you won't see much activity. As noted earlier, these are often codeless device drivers only to get them installed early for use by other device drivers. In the case of AM5, in fact, Microsoft has a non-device-driver version of HIMEM.SYS called AM5MMB.LOAD to use the Windows 3.1 SETUP when the user doesn't have HIMEM.SYS or another AMS driver installed. For some of these drivers, *do* deal with device-driver packets, particularly generic IOCTL calls. For example, 386MAXND and 386MAXS have generic IOCTL calls that



are used by Windows for these, on the other hand it is probably easier to watch calls to INT 21h AH=44h rather than inspect the lower-level device driver packets.

### Watching XMS

While intrinspy also does direct entry points account for many of the communications between software and the PC, it is who means all. To use XMS for example, a program makes an indirect call to INT 2Fh AX=4310h and the address returned in EBX is then used for a subsequent call to XMS services. Simply intercepting INT 2Fh and watching for AX=4310h will not get you very far.

XMS also can be intercepted in an INTRSPY script using another form of the INTRUPT command. In this case, the interrupt number or device driver name is replaced by the special keyword XMS, as seen in the XMSWATCH.SCR script in Listing 5-4.

### Listing 5-4: XMSWATCH.SCR

```

, XMS.SCR
structure move fields
  len (dword,hex)
  src (word,hex)
  src_ofs (dword,hex)
  dest (word,hex)
  dest_ofs (dword,hex)

intercept xms
function 0 on_entry output XMS 00 - get xms version number
function 1 on_entry output XMS 01 - request HMA " dx
function 2 on_entry output XMS 02 - release HMA
function 3 on_entry output XMS 03 - global enable A20*
function 4 on_entry output XMS 04 - global disable A20*
function 5 on_entry output XMS 05 - local enable A20
function 6 on_entry output XMS 06 - local disable A20*
function 7 on_entry inc #1
  ; Query A20 state. Too many of these!
function 8 on_entry output XMS 08 - query free extended memory
  on_exit same line -> dx
function 9 on_entry output XMS 09 - allocate extended memory dx
function 0Ah on_entry output XMS 0A - free extended memory dx
function 0Bh on_entry output XMS 0B - move extended memory
  ; dx si ->move len
  ; from (dx si ->move src) / / (dx si ->move src_ofs)
  ; to (dx si ->move dest) / / (dx si ->move dest_ofs)
function 0Ch on_entry output "XMS 0C - lock extended memory dx
function 0Dh on_entry output XMS 0D - unlock extended memory dx
function 0Eh on_entry output XMS 0E - get handle info
function 0Fh on_entry
  output "XMS 0F - realloc extended memory dx " dx " bx
function 10h on_entry output XMS 10 - request UMB dx
function 11h on_entry output "XMS 11 - release UMB dx
function 12h on_entry output "XMS 12 - realloc UMB

include "prog"

```

To keep it in line with Listing 5-4, this includes the XMS version 2.0 specification calls. With a little more effort you could include additional calls or 3.0 calls.

The last line of XMSWATCH.SCR shows another INTRSPY feature. Common scripts now be included using the INCLUDE statement with the use of a single #include m/c or #include m/c by Pascal. The PROG.SCR file shown in Listing 5-5 records the start of program execution. Later in the chapter, a file executor called EXEC.SCR provides an INTRSPY machine model. In its PROG.SCR component, it is assumed that you will include in many of your scripts a call to a file called exactly what proportion of the criteria your primary script is recording actually relates to a particular program.

**Listing 5.5: PROG SCR**

```

; PROG SCR
Intercept 21h function 4bh on_entry
    output "-----"
    output (ds:dx->byte,ascii,64)

```

When AMISWATCH while Windows is starting up, the AMISWATCH results show, for example, that Windows gains a free extended memory. It calls AMS function 8 to query the size of the largest block of available extended memory, and then calls AMS function 9 to allocate that largest available block.

Intercepting a particular function call via AMS function 7 (Query A20 State) a counter is used to process calls. Windows typically calls this function thousands of times in a very brief period. To get precise timing, one of each of these would produce a very minor error report.

Most of the AMISWATCH report will be devoted to AMS function 0Bh Move. In the normal course of operation, it moves 4000 bytes, some 950 of which were function 0Bh. It has a fixed source and target, except for using the workhorse functions of the AMS interface. AMISWATCH can intercept this very normal call with the source and destination output as xxxx yyyyzzzz. If xxxx is zero, it represents conventional memory address yyyy-zzzz. Otherwise, yyyy is an extended memory address and zzzz is 32 bits of the into the extended memory block. EMM.

Intercepting the interrupt traps the AMS alerting for a brief period during the Windows load process. Because Windows takes over responsibility for AMS and because as a result it does some amount of critical work, INTRSPY must temporarily lose track itself while Windows makes the transition to 386 Enhanced mode.

INTRSPY knows that Windows is starting up through the useful Int 21h function 16h All pointed to Windows. Among other things, this utility provides hooks to allow INTRSPY to not only intercept starting and stopping of Windows, but to perform any necessary preparations and even Windows to load stored data, if necessary. The MUI-1 module that will be presented in Chapter 9 uses this hook to store a series of capabilities between MUI-1 and Windows.

Normally, Windows will have Windows starting code with a 386 interrupt mode, both a standard and a standard. Such may be 0Bh in standard mode and in memory INRs that the Windows interrupt is a patch. A second pair of calls subtractions 09h and 06h to get INRs that interrupt mode is a patch that Windows is in standard mode is complete, respectively. INTRSPY disables AMS work during startup and concludes to allow the V86MMKR module of Enhanced mode Windows to intercept ARPI invocation at the AMS endpoint. This instruction will generate a new set of calls. Windows traps to receive control. INTRSPY then reinstalls its own patch, all on top of the intercept AMS only before passing the word to Windows.

**Dynamic Hooks**

INTRSPY can provide a mechanism for watching calls to any function, not just AMS. In fact, AMS can be a source for this particular identification hook. You can specify your own hooks for other addresses that are not intercepted. This feature is very powerful and allows the trapping of callbacks.

Using the INTRSPY HOOK facility, for example, a script can be written to watch the completion of NetBIOS requests. NetBIOS can appear to communicate asynchronously over the network, but is a synchronous protocol in receiving some. NetBIOS goes about the business of sending or receiving data. One way to accomplish this is by specifying a post routine as one of the parameter fields at the NetBIOS endpoint hook. NetBIOS is passed to Interrupt 21h when requesting a NetBIOS network communication service. At the completion of whatever task has been requested, NetBIOS calls the post routine, notifying the application.

The most significant difference between previous INTRCEPI commands and the entry point identifier is that the address that is to be intercepted is supplied, not at compile time, but as a

result of another event. This is achieved using a HOOK statement with the same identifier, as shown in Listing 5-6 (NETBIOS.SCR).

### Listing 5-6: NETBIOS.SCR

```
; NETBIOS.SCR
structure ncb fields
    cmd          (byte)
    rc           (byte)
    loc_sessno   (byte)
    name_no      (byte)
    buffer       (dword,ptr)
    buflen       (word,dec)
    dest_name    (byte,char,16)
    srce_name    (byte,char,16)
    rto          (byte,dec)
    slo          (byte,dec)
    post_routine (dword,ptr)
    adap_no     (byte)
    cmd_rc      (byte)
    reserved     (byte,hex,14)

Intercept post_routine
on_entry
    output
    output "-----"
    output "Post routine triggered for NCB at 'es': 'bx"
    output " = completed" (es:bx->ncb.cmd_rc)

Intercept $ch
on_entry
    output
    output "-----"
    output "NCB at es: bx submitted" (es:bx->ncb)
    if ((es:bx->ncb.cmd > 127)
        output "Command is NNAME!"
    if ((es:bx->ncb.cmd > 127) EE ((es:bx->ncb.post_routine) != 0)
        hook post_routine = (es:bx->ncb.post_routine)
        newline " Post routine used"
on_exit
    output "On exit AL is: a
```

The first INTERCEPT statement refers to an identifier, `post_routine`. When the script is compiled, there is no address associated with `post_routine` since the NCB whose `post_routine` we want to trap, has not yet been created. However, our second INTERCEPT statement is watching NetBIOS intercepts INI \$ch. We can use a pointer to intercept NetBIOS commands when the NCB is for processing; the intercept can determine if a post routine has been specified (that is, the `POST` field in the supplied NCB structure is not NULL). If a post routine has been specified, the HOOK statement uses the address in the `POST` field as the address to be patched by the post routine intercept.

INSTRIFY's dynamic treatment of intercepts (intercepts on the one hand and entry point identifiers on the other) can be identified to the difference between static and dynamic linking. Static linking establishes the addresses for all various calls between modules when a program is built. Dynamic linking only establishes those intercept module call addresses at run time. In the same way, INSTRIFY asks intercept and device driver INTERCEPT statements to the appropriate addresses at compile time, but links a entry point identifier to the appropriate address when the associated HOOK statement within another INTERCEPT is invoked.

In Listing 5-6, notice how an entire NCB structure resulted from the single statement `ES:BX=NCB`. When we asked INSTRIFY to take care of displaying every field in a structure, including the field's name, the resulting output can be seen in Figure 5-6, which shows output from NETBIOS.SCR for a single NetBIOS command. As can be seen from the NCB structure, the

command was `22h` (this is a Send Broadcast Datagram `22h` with the no wait bit `80h` set). The NCB was located at `0CDD00DA`. The actual data to send was located at `0CDD0048` (NCB.BUFFER = 18 decimal bytes were to be sent: NCB.BUFFER\_N). We could, of course, have displayed the contents of this buffer as well.

**Figure 5-6. Output from NETBIOS.SCR for a NetBIOS Command**

```

NCB at 0CDD:00DA submitted
NCB.CMD                42
NCB.RC                 1: 00
NCB.LOC_SESSNO        1: 00
NCB.NAME_NO          - 02
NCB.BUFFER            0CDD:0048
NCB.BUFFER_N         18
NCB.DEST_NAME
NCB.SRC_NAME
NCB.RTO               1: 0
NCB.STO               1: 0
NCB.POST_ROUTINE     1: 066B:003B
NCB.ADAP_NO           00
NCB.CMD_RC           1: 00
NCB.RESERVED         00 00 00 00 00 00 00 00 00 00 00 00
00 00 00
Command is NOMWAIT - Post routine used
-----
Post routine triggered for NCB at 0CDD:00DA - completed 00
On exit AL is 00

```

Since the goal of this example is to illustrate the `INTRSPY HOOK` statement, the important thing in Figure 5-6 is the end, where the post routine was called at some later point for the initial no wait asynchronous send broadcast datagram command. As you can see, the completion code (NCB.CMD\_RC) was a zero, indicating success.

The above is a crude example of the `HOOK` function's use in watching callback functions, not `EXITS`. The critical caveat associated with its use in at least the NetBIOS context. The problem is always the last result when the program ends. Windows enhanced mode. Since an application running in enhanced mode may set the no-wait mode at the precise moment that its NetBIOS event completes, a virtual delay before `WIN32API` is installed to undertake anything other than the virtual memory management and task switching necessary to switch since the VM and the callback. In order to do this is the `INTR` hook at the application's expense, since the NCB is intercepted by the VxD and the post routine address is replaced with the protected mode address of a thunk that will be invoked before the post routine is called. Additionally, the buffer address is converted to the address of a protected mode holding buffer, and the real mode callback can be called.

The above should be a further warning that `INTRSPY` is running in the same DOS window and to a large extent purposes replaces the application temporarily. Using the `ON_ENTRY` class's window exit picking any addresses before they are installed. Unfortunately, not Windows gets there first, using a `ARPL` instruction `VM86` breakpoint, and passes control back to the VM's real mode `crash` handler again. Thus, even in the `ON_ENTRY`'s case, we are too late. This all means that in a DOS box, breaking a NetBIOS callback will not lead to the desired results and will in all likelihood crash the VM.

Note the similarity here to cover earlier problems with `VMS`. For a program like `INTRSPY`, Windows enhanced mode is a tough application with which to be compatible.

## INTRSPY User's Guide

INTRSPY is actually two programs: INTRSPY.EXE and CMDSPY.EXE, whose operating instructions follow.

### Using INTRSPY.EXE

The following command is used to run INTRSPY.EXE:

```
INTRSPY [-rnnnn] [+*]
```

where *rnnnn* specifies the amount of memory INTRSPY is to allocate for compiled script and result storage (default is 8K). For example, the following command runs INTRSPY with an allocation of 24,000 bytes of code and results space:

```
C:\VMD002\CHAP5>INTRSPY r24000
```

### Using CMDSPY.EXE

The following command is used to run CMDSPY.EXE:

```
CMDSPY [COMPILE [d ][path]inputfile[.ext] [param-1 [param-2 ...]]
[REPORT]
[RESTART]
[FLUSH]
[STOP]
[UNLOAD]
```

where

- COMPILE** compiles a script and instructs INTRSPY to begin monitoring interrupts and entry points as well as storing results specified in *inputfile[.ext]*, which contains script to execute in the user-defined delay. Any errors in the script stop it, and the results area is indicated if the extension is omitted. SCR script is assumed. For example, the following DOS command line compiles the script TEST.SCR:

```
C:\VMD002\CHAP5>CMDSPY COMPILE TEST
```

- REPORT** instructs INTRSPY to report the results accumulated since the script was compiled or since the last time a CMDSPY FLUSH was issued (see below). CMDSPY formats the results in the results buffer, as specified in the current script, and writes them to STDERR. The display can be collected on file. For example:

```
C:\VMD002\CHAP5>CMDSPY REPORT > TEST.LOG
```

- STOP** instructs INTRSPY to stop monitoring interrupts and entry points, but to preserve the results area.
- RESTART** instructs INTRSPY to restart monitoring interrupts and entry points after a STOP or RESTART command, on the basis of the currently compiled script.
- FLUSH** instructs INTRSPY to clear the results area, but to save the current script active.
- UNLOAD** causes the INTRSPY ISR to attempt to unload itself from memory. A relative script is stopped. There are circumstances under which INTRSPY cannot be unloaded.

### Script Language

The script language allows the main constructs:

- INCLUDE**, which includes another input file.
- STRUCTURE**, which defines a data structure.
- SCRIPTLINE**, which specifies an interrupt number, device, driver, strategy, a script routine, or an entry point identifier to be used with a **HOOK** statement. It also specifies optional

function and substitution numbers or ranges, together with the entry and exit processing to be done when the intercept is triggered.

- **RUN** which allows a DOS program to be EXEC'd from within the script and a built-in debugger to be used within intercepts.
- **GETRVAL** which allows intercepts to be generated, usually to obtain an address to be associated with an entry point identifier.
- **REPORT**, **FLUSH STOP**, and **RESTART** all of which work exactly like their command-line counterparts.

### Syntax

A script file is a **SCRIPT**. All white space is ignored, except within literal strings used for results or point identifiers. Comments and multiple lines may be used for readability or lack thereof, as the user may find it confusing to copy a file to debug programs which begin with a semicolon anywhere on a line. This is particularly true for **INTRSPY** which can appear anywhere in a script and are replaced by the DOS command line. The following is thus a valid **INTRSPY** script:

```
; INTERCEPT.SCR
intercept X1
function X2
  X3 X4 X5 X6 X7 X8 X9
```

This script could also be used to intercept operations that didn't describe their own separate scripts, for example:

```
C >cmdspy compile intercept 21h 52h on_exit output es : ' h'
```

The simplest possible valid **INTRSPY** script is thus:

```
; SCRIPT.SCR
X1 X2 X3 X4 X5 X6 X7 X8 X9
```

This request in the command script be placed on the DOS command line:

```
C >cmdspy compile script intercept 21h function 52h on_exit output es : ' h'
```

### INCLUDE Syntax

The following syntax is used for **INCLUDE** statements:

```
INCLUDE [d:]path[intrspy]of ext1 [param-1] [param-2] ...
```

This syntax includes the specified file and substitutes **param-1**, **param-2**, and so forth in the source for the strings **%1**, **%2**, ..., **%N** respectively, which found if the extension is omitted. **SCR** is assumed. Some examples:

```
include "1" ; include script named on command line
include "foo" ; include foo.scr
include "foo X1" ; include foo.scr, passing arg from cmd line
include "foo bar" ; include foo.scr, passing arg "bar"
```

### STRUCTURE Syntax

Some example structures, such as **param blk** and **blk** in Listing 5.2 (**FDOPEN SCR**) and **reqn** in Listing 5.3 (**ADD SCR**), were shown earlier.

To formalize some **STRUCTURE**, first define a *field definition* as:

```
field-type [,field-dtyp-type [,field-dup]]
```

where:

- *field-type* can be **BYTE**, **WORD**, **DWORD**, or a preprocessor defined structure name

- *field disp type* is `HEX`, `BIN`, `DEC`, `INT`, `PIR`, `CHAR`, `PASCAL` a string with the count as its first byte, `ASCII` a zero-terminated string, `DUMP` a combination of `HEX` and `ASCII`, `STRUCT` if field type was a structure name.
- *field disp* is the number of elements if the field should be treated as an array or the length of the field if, for example, a string. In the case of a field being defined with the structure definition, the field disp may be a numeric literal or one of the predefined constants. In the case of a definition with no output element (see `INTRCEPT` syntax), field disp may also refer to a register or a constant. Then `STRUCTURE` syntax looks like this:

```
STRUCTURE struct name FIELD$
  field-name1 ((field-definition))
  [field-name2 ((field-definition))]
  ...
  [field-nameN ((field-definition))]
```

Note that struct name must be a unique structure identifier and that field name must be unique within struct name. There are thirty-eight characters for both struct name and field name, and indeed for any INTRSPY words. This number should be adequate for almost any application.

### INTERCEPT Syntax

Main examples of `INTRCEPT` statements were presented in the first portion of this chapter, but as is now formally below `INTRCEPT`. To begin with, let us define an intercept element as an interrupt number, a device driver string, an interrupt routine, or an entry point identifier. For example,

```
INTERCEPT 21h ; intercept an interrupt
INTERCEPT STRATEGY( foo ) ; intercept device driver Strategy routine
INTERCEPT INTERRUPT( foo ) ; intercept device driver Interrupt routine
INTERCEPT hook_proc ; intercept entry point identifier (see below)
```

Then, define an *output element* as

```
REGS or
reg_or_flag or
(reg_or_flag, field disp type) or
#counter or
(#counter, field-disp-type) or
"string literal" or
(segval+seg_adjust offset[additional-offset])->memory-reference) or
predefined-constant or
(predefined-constant, field-disp-type)
```

where

- *reg or flag* is a name of `REGS` or `FLAG` register or a flag.
- a *flag* is defined by `API.W`, where `x` may be one of the following: `D` (direction), `I` (interrupt), `trap`, `S`, `z`, `o`, `ov`, `V`, `cx`, `ov`, `P` (parity), `ac` (carry). For example, `CF` (API.W) or `API.W`.
- *#counter* is one of 16 bits in 32-bit count registers identified by #1 through #16.
- *reg or flag* is a register or a flag constant or `REGS` or `FLAG` constant.
- *seg adjust* is an optional numeric constant to be added to or subtracted from the segment value. For example, `es 10 +mbch`.
- *offset* is an optional numeric constant or built-in constant.
- *add or sub* is an optional numeric literal increment or decrement. For example, `(es:bx) [2] sub 31`.
- *memory reference* is one of the following:

o `BI`, `WORD`, `DWORD` field disp type field disp, both defined in `STRUCTURE` above.

- STRIKE
- STRUCT FIELD

\* STRUCT dword field: memory reference, where dword field means that the field to be used as a pointer must have been defined as DWORD struct and field must be predefined

- predefined constant is one of those defined in the section "Predefined Constants", below

Define an *action statement* as one of the following:

- OI IPU [ output element ] output element [ output element ] which starts on a new line
- SAMI IPI [ output element ] output element [ output element ]], which attempts to append elements to an existing line of output
- SRI VM r, g, s, name, which outputs raw VM II characters from the 8 bit register specified
- DEB G, which, if the intercept occurs during a RUN statement, enters a pop-up debugger
- INC R #, counter, which increments the specified 32 bit count register
- DEC R #, counter, which decrements the specified 32 bit count register
- ZI R #, counter, which resets the specified 32 bit count register to zero
- HOOK entry point identifier hook-spec, where `entry` may be replaced by the word 'IO', and hook-spec may be register or signal; seg-adjust, ofval, additional offset, memory-ref, rler, vler, where memory reference may not include a field disp type or field disp

Define a *test clause* as

(operand operator operand)

where operand may be a

reg\_or\_flag or  
#counter or  
(segval+/- seg-adjust, ofval [additional offset]) >memory-reference) or  
predefined-constant

where memory reference is one of the following:

- BYTE WORD DWORD
- struct field where field must be defined as a non-dup BYTE WORD DWORD
- struct dword field: memory reference, where dword field means that the field to be used as a pointer must have been defined as DWORD struct and field must be predefined

and where operator may be `==` equal, `!=` not equal, `>` greater than, `>=` greater than or equal to, `<` less than, or `<=` less than or equal to

Next define an *if clause* as follows:

IF test-clause [AND test-clause [OR test-clause [ ]]]  
[action-statements]

where &Z may be used interchangeably with AND and &V may be used interchangeably with OR

Use the two keywords ON ENTRY and ON EXIT to describe pre- and post-processing of an intercept

Finally, the two additional keywords FUNCTION and SUBFUNCTION provide a shorthand and set documenting form of IF AH == nn and IF AL == nn. However, those tests would likely produce incorrect results at the post-processing stage, since AX may have been modified, so the FUNCTION and SUBFUNCTION keywords recall the values that AH and AL had on entry to the intercept

FUNCTION ah-value  
SUBFUNCTION al-value

Then, INTERRUPT syntax looks like this



```

INTERCEPT [intercept-element]
  [action-statements]
  [test-clauses]
  [FUNCTION fnctn-number [fnctn-number ...]
    [action-statements]
    [test-clauses]
  [SUBFUNCTION sfnctn-number [sfnctn-number ...]
    [action-statements]
    [test-clauses]
    [ON_ENTRY
      [action-statements]
      [test-clauses]]
    [ON_EXIT
      [action-statements]
      [test-clauses]]]
  [SUBFUNCTION sfnctn-number [sfnctn-number ...]
    :
    ]
[FUNCTION fnctn-number [fnctn-number ...]
  :
  :
  ]

```

Although the ON\_ENTRY and ON\_EXIT keywords are shown within the scope of the SUBFUNCTION as if in the implied block hierarchy, they need not be; they could have been shown at a higher level below INTERCEPT with FUNCTION and SUBFUNCTION as lower levels in that hierarchy. For example:

```

INTERCEPT intercept-element
  ON_ENTRY
    FUNCTION fnctn-number
    .. etc ..

```

### GENERATE Syntax

The GENERATE construct allows an interrupt to be generated from within a script. This allows scripts to be self-sufficient (see INTERRUPT.SCR, Listing 5.8). Without it, some other program must be run to generate the wanted interrupt.

The GENERATE construct uses some of the INTERCEPT syntax. There is an ON\_ENTRY section and/or an ON\_EXIT. The difference is that in GENERATE one or the other section is required. The ON\_ENTRY clause allows for the setting up of registers for the interrupt, and the ON\_EXIT clause, if it allows, is used to only one statement, the HOOK command, external to the INTERCEPT syntax.

The syntax of the ON\_ENTRY clause, where hook-spec is as defined in the INTERCEPT section, is simple:

```

ON_ENTRY
  reg = value [, reg = value [, reg = value]]

```

The syntax of the ON\_EXIT clause is

```

ON_EXIT
  HOOK entry-point-identifier = hook-spec
  [HOOK entry-point-identifier = hook-spec]
  ...

```

### RUN Syntax

The RUN syntax looks like this:

```

RUN "[drive:]path]program[.ext] [parm1 [parm2 ...]]"

```



**Figure 5-7: The INTRSPY Debugger**

```

#####
mm          INTRSPY DEBUGGER
#####
AX  BX  CX  DX  BP  SI  DI  DS  ES  SS  SP      CS  IP      FLAGS
3000 377E 0000 0437 3094 3330 0BE0 4173 4173 4173 308E 3C34-3009  0d162aPc
>>d ds bx->byte,asc"1,04
C:\MSC\BIN\c1_err
>>>

```

Note that at the debugger prompt `>>>`, you can type an expression similar to those enclosed in INTRSPY scripts. As noted earlier, the `DEBUI` statement only invokes the debugger when used in a script with a `RUN` statement, without `RUN DEBUI` as a NOP.

### Predefined Constants

The following constants are provided within the script language and the debugger:

- `OS MAJOR` and `OS MINOR`: The major and minor DOS version numbers obtained from INT 21h Function 30h GetVersion.
- `100`, `SEG`, and `OFFS`: The segment and offset portions of the EBP/EIP of last SysVars, obtained from INT 21h Function 52h SysVars.
- `SDA SEG` and `SDA OFFS`: The segment and offset portions of the address of the primary Swapable DOS Area obtained from INT 21h Function 51006h Get Swapable DOS Area. This is useful in specifying the network redirector and for queries to DOS such as the current PSP. For example,

```

Intercept 21h
function 50h
on_entry
  if (bx != (sda seg sda ofs[10h]->unrd))
    output "PSP changed to" bx

```

### Error Messages

**CMDSPY Compilation Messages** CMDSPY generates explicit script compiler error and warning messages as it detects a problem, like the following shows sample CMDSPY output when an error is encountered:

```

File: BAD.SCR
Line 2:   field1      (word, ptr, 3)
          ***

```

Error: PTR format can only be used with DWORD fields

The first line shows that `PADMSR` was the last CMDSPY pass, so it is known that the compiler had the error. The next line shows that it occurred twice. The `***` and the asterisks show the location that was in error. The `***` are the error messages. The following shows one of the warning messages that shows only a potential problem but that draws attention to something that may be important and that may affect the results that you expect:

```

File: DD_INTR.SCR
Line 3: Intercept interrupt("A")
          ***

```

Warning: Same block device driver supports multiple drives  
16 bytes of code generated for 'A' driver interrupt script

In this example, the compiler has found that the same `device` driver handles more than A, presumably B and C, as well. The results that you obtain will reflect not only A, but also the other drives. This condition is not a error, so compilation proceeds.

**CMDSPY and INTRSPY in Operation** Both CMDSPY and INTRSPY respond to a `^C` or `^P` command and display with a summary of their command-line options. In addition, they generate various operator messages—many of which are quite informative, but some of which are errors. All errors, by definition, are fatal, so I will only deal with some of the more interesting messages here.

One message you'll encounter frequently if you are watching INTRSPY and use `STOP` or `UNLOAD` to terminate the script is

```
Cannot disable current INTRSPY int 21h handler.
It is handling the DOS Exec of CMDSPY.
Try again, need delay and it will disable successfully?
```

When INTRSPY watches INTRSPY, INTRSPY calls `INT21H` handlers for those functions that are not the subject of the script you are watching. I have arranged the processing for the script as simply passed on. However, if you are watching the execution of DOS, your `INT21H` handler is bypassed. Execution is back through the INTRSPY INTRSPY handler. Since CMDSPY uses `STOP` by the DOS command interpreter issuing a `^C` or `^P`, the `INT21H` handler cannot be reinstalled until that function has returned. The handler cannot catch the program from `INT21H` so that subsequent INTRSPY calls do not pass through it. I have arranged for the CMDSPY `STOP` or `UNLOAD` command to be issued, the handler can be reinstalled successfully.

The following message is displayed if you are watching a function that is not in the

```
Cannot disable current INTRSPY script.
Receive any subsequently loaded TSDs and try again.
```

The error message appears because the TSDs are included in the accompanying disk as if to install a PHANTOM program as discussed in chapter 8. Then you install PHANTOM. Finally, attempt to simulate a function that is not in the TSDs. The "Cannot disable" message is generated because the handler for the `INT21H` function that was installed before the TSD handler was installed, CMDSPY, was able to re-install itself to catch script stop points at our handler. If it does not—as is the case with the PHANTOM installation, it is installed before the start to remove the TSDs and then the program can catch functions. PHANTOM is then forced to keep up its responsibility to prior INTRSPY handlers. I cannot set the previous vector, an address in our resident code, and will pass `INT21H` calls to the DOS INTRSPY handler. Since we do not load the contents of memory at the address that PHANTOM claims to be, it would be indeterminate.

The `INT21H` handler will be able to re-install itself to catch script stop points. At this point, we are unable to catch the `INT21H` handler command and catch it. Instead, it will be installed. So, in fact, you have to be careful to avoid other INTRSPY that have installed after INTRSPY.

## INTRSPY Utility Scripts

The following look at some of the applications that use the INTRSPY language.

### UNDOC

In figure 5.8, as the `UNDOC` INTRSPY package shows which undocumented DOS calls a program makes. Note that the program is presented in chapter 1. In fact, we were interested only in the initial `INT21H` INTRSPY call. The `UNDOC` script is using two pieces of system software: the Microsoft CD-ROM Executive, MSCDEX, and the NetWare shell, NET5. A step-by-step account is given in figure 5.8. A screenshot shows some CMDSPY output, which we have, until now, been omitting as extraneous.

**Figure 5.8. INTRSPY UNDOC.SCR Watching MSCDEX and NET5**

```
C:\BIN> intspy                               (1)
8192 bytes allocated for code and results.
```

INTRSPY v2.00 loaded

C:\BIN> cmdspy compile undoc (2)

715 bytes of code generated for Int 21h script  
 65425 bytes of code generated for Int 20h script  
 65425 bytes of code generated for Int 27h script  
 27 bytes of code generated for Int 2Eh script  
 UNDOC.SCR compiled successfully.  
 2928 bytes of space used for code and overhead  
 5264 bytes of results space available.  
 Handler for Int 21h started.  
 Handler for Int 20h started.  
 Handler for Int 27h started.  
 Handler for Int 2Eh started.

C:\BIN> mscdex /d:\xrw1001 /i g (3)

MSCDEX Version 2.20  
 Copyright (C) Microsoft Corp. 1986, 1987, 1988, 1989, 1990 All rights reserved  
 Drive G: = Driver XRW1001 unit 0

C:\BIN> cmdspy report (4)

Handler for Int 21h stopped  
 Handler for Int 20h stopped  
 Handler for Int 27h stopped.  
 Handler for Int 2Eh stopped.  
 UNDOC.SCR compiled successfully.

----- -- Start of Report -----

C:\BIN\MSCDEX EXE (5)

2152: Get List of Lists: 011C 0026

----- TSR ----- (6)

C:\BIN\CMDSPY EXE  
 2152: Get List of Lists: 011C 0026

----- - - End of Report -----

All counters zero

148 bytes of results reported

Intrspy returned successful status

C:\BIN> cmdspy flush (7)

INTRSPY Buffer flushed

C:\BIN> net5 (8)

.

C:\BIN> cmdspy report (9)

Handler for Int 21h stopped  
 Handler for Int 20h stopped  
 Handler for Int 27h stopped.  
 Handler for Int 2Eh stopped.  
 UNDOC.SCR compiled successfully.

----- Start of Report -----

C:\NOVELL\NET5.COM

2134 InDOS flag 0116 0321

2134 InDOS flag 0116 0321

2152 Get List of Lists 0116 0026

2150 Set PSP 0486

2150 Set PSP 0B5B

----- TSR -----

2151 Get PSP 0486

2152 Get List of Lists 0116.0026

```

C:\BIN\CMDSPY.EXE
2152: Get List of Lists: 0116.0026

```

```
----- End of Report -----
```

All counters zero.

148 Bytes of results reported.

Intrapy returned successful status.

At 1. INTRSPY loads using the default allocations for interrupt handling code and results space. At 2. UNDOC.MKR is completed. The version of UNDOC.MKR shown in Listing 5 includes some functions, such as INI 21h, AII 34h and AII 90h, which were documented for DOS 5.0.

### Listing 5.7: UNDOC.SCR (Includes Some Documented Calls)

```

** UNDOC.SCR
Intercept 21h
  function 1fh on_exit output "211f Get Default DPM " DS " BX
  function 32h on_entry output "2132 Get DPM " DL
  function 34h on_exit output "2134 InDOS flag " ES " BX
  function 50h on_entry output "2150 Set PSP BX
  function 51h on_exit output "2151 Get PSP BX
  function 52h on_exit output "2152 Get List of Lists " ES " BX
  function 53h on_exit output "2153 Translate BPP"
  function 5dh subfunction 0bh
    on_exit output "21500b Get BOSSMAP DS " SI
  function 60h
    on_entry output "2160 Canon file. (DS SI->byte,asciiz,64)
    on_exit sameline =>> (ES DI->byte,asciiz,64)
  function 25h
    on_entry
      if not 28h) output "Select INT 28h E80 busy loop"
**
** Use the next functions and ints 20h and 27h to show which
** program made the undocumented DOS calls, and to show termination
**
  function 4bh
    subfunction 0bh
      on_entry
        output (DS DI->byte,asciiz,64)
    subfunction 01h
      on_entry
        output "214801 E8EC debug (DS DI->byte,asciiz,64)
  function 4ch on_entry output "-----
  function 35h on_entry output "----- TSR -----"
Intercept 20h on_entry output "-----"
Intercept 27h on_entry output "----- TSR -----"
Intercept 2eh on_entry output "2E Execute command"

```

At 3. MMSDFX runs. At 4. a screen report of the results to this point is generated. So that it is possible to distinguish which calls were generated by which program in the event that we want to analyze events for a while, or in the event that one program spawns others, the script monitors the DOS INI and termination functions and interrupts. The line of hyphens at 5 is the end of the run of CMDSPY from 2, that loaded the script. The line following it shows that MMSDFX has started. At 6. MMSDFX is terminated but stayed resident. In the next line, we see the invocation of CMDSPY at the next line 4. At 7, the results space is cleared; this action isn't really necessary, but it ensures that the next report is uncluttered to what happened on INI 5, without reiterating the MMSDFX results. At 8. INI 5 starts, and at 9, the results of its loading are reported.

In this experiment, we see that MMSDFX calls INI 21h function 52h at load time, and that INI 5, in addition to function 52h, also calls functions 34h and 50h. The reasons these programs

make these particular DOS calls should be clear to you from other chapters. The DOS List of Lists retrieved with function 52h contains a pointer to the DOS 4 current Directory Structure. MSCDEX alters the CDS. Functions 34h and 50h are both important for ISRs like NE15.

## LISTOFLST

In Chapter 2, we went through a fairly laborious process using C to display SysVars, the DOS List of Lists. It is simpler to do this with INTRSPY, as seen in Listing 5-8 (LISTOFLST.SCR).

### Listing 5-8. LISTOFLST.SCR Displays SysVars

```

; LISTOFLST.SCR
; INTRSPY script to examine DOS List Of Lists (INT 21h Function 52h)
structure list_20 fields ; DOS 2.x
  share_retry_count (word, dec)
  retry_delay (word, dec)
  curr_disk_buff (dword, ptr)
  unread_con (word, dec)
  mcb (word)
  dpb (dword, ptr)
  file_tbi (dword, ptr)
  cclock (dword, ptr)
  con (dword, ptr)
  num_drives (byte, dec)
  max_bytes (word, dec)
  first_disk_buff (dword, ptr)
  nul (byte, dump, 18)
structure list_30 fields ; DOS 3.0
  share_retry_count (word, dec)
  retry_delay (word, dec)
  curr_disk_buff (dword, ptr)
  unread_con (word, dec)
  mcb (word)
  dpb (dword, ptr)
  file_tbi (dword, ptr)
  cclock (dword, ptr)
  con (dword, ptr)
  num_blk_dev (byte, dec)
  max_bytes (word, dec)
  first_disk_buff (dword, ptr)
  curr_dir (dword, ptr)
  lastdrive (byte, dec)
  string_area (dword, ptr)
  size_string_area (word, dec)
  fcb_tbi (dword, ptr)
  fcb_y (word, dec)
  nul (byte, dump, 18)
structure list_31 fields ; DOS 3.1+
  share_retry_count (word, dec)
  retry_delay (word, dec)
  curr_disk_buff (dword, ptr)
  unread_con (word, dec)
  mcb (word)
  dpb (dword, ptr)
  file_tbi (dword, ptr)
  cclock (dword, ptr)
  con (dword, ptr)
  max_bytes (word, dec)
  disk_buff (dword, ptr)
  curr_dir (dword, ptr)
  fcb (dword, ptr)
  num_prot_fcb (word, dec)
  num_blk_dev (byte, dec)

```

```

lastdrive (byte, dec)
nul (byte,dump,18)
num_join (word, dec)

Intercept 21h
on_exit
  function 52h
    output "DOS Version is " (OS_MAJOR, dec) "." OS_MINOR
    if (OS_MAJOR == 2)
      output (es:bx[-12]->list_20)
      if (OS_MAJOR == 3) and (OS_MINOR == 0)
        output (es:bx[-12]->list_30)
      if (OS_MAJOR == 3) and (OS_MINOR != 0)
        output (es:bx[-12]->list_31)
        output ""
        output "CON device header"
        (es:bx[-12]->list_31.con->byte,dump,18)
        output ""
        output "CLOCK device header"
        (es:bx[-12]->list_31.clock->byte,dump,18)
      if (OS_MAJOR > 3)
        output (es:bx[-12]->list_31)
        output ""
        output "CON device header"
        (es:bx[-12]->list_31.con->byte,dump,18)
        output ""
        output "CLOCK device header"
        (es:bx[-12]->list_31.clock->byte,dump,18)
  function 52h
    output "
j, Issue an Int 21h/AH=52h
generate 21h
on_entry
ah = 52h

report

```

Note that `output` displays the entire structure with one `OR IPU 1` statement. `CMIDSPY` takes care of formatting output according to the format options specified in the structure itself. When you are displaying an entire structure with a single `OR IPU 1` statement, `CMIDSPY` also takes care of displaying the field names. For example, sample output from `INTRSPY SYSVAR` is shown in Figure 5-9. This output was produced under DOS 5.0; the structure is called `LIST_31` because the fields we are examining here haven't changed since `1005 3 1`.

**Figure 5-9: INTRSPY Displays SysVars (the List of Lists)**

```

DOS Version is 5.00

LIST_31.SHARE_RETRY_COUNT      : 3
LIST_31.RETRY_DELAY            : 1
LIST_31.CURR_DISK_BUFF        : FFFF-A078
LIST_31.UNREAD_CON             : 0
LIST_31.MCB                    : 0251
LIST_31.DPB                    : 0116 1362
LIST_31.FILE_TBL               : 0116 00EC
LIST_31.CLOCK                  : 0070 0059
LIST_31.CON                    : 0070 0023
LIST_31.MAX_BYTES              : 512
LIST_31.DISK_BUFF              : 0116 0060
LIST_31.CURR_DIR               : 0381 0000
LIST_31.FCB                    : 034F 0000
LIST_31.NUM_PROT_FCB           : 0
LIST_31.NUM_BLK_DEV            : 3
LIST_31.LASTDRIVE              : 12

```



```

LIST_31.MUL
0000 | 00 00 9b 02 04 80 c6 0b cc 0b 4e 55 4c 20 20 20 | . MUL |
0010 | 20 20
LIST_31.NUM_JOIN : 0

CON device header
0000 | 35 00 70 00 13 80 f5 06 00 07 43 4f 4e 20 20 20 | 15.p . CON |
0010 | 20 20

CLOCK device header
0000 | 68 00 70 00 08 80 f5 06 39 07 43 4c 4f 43 4b 24 | 1x.p 9.CLOCKS|
0010 | 20 20

```

The `CON` or `STRM` command now for this version allows us to generate an output so that we can trigger a record of the contents of `hexVars` without waiting for some program to first call `INT 21h AH=52h` for us. Note, also, how a hex dump can be produced with the `DD.M` output code if for example `DD IPI 1 conv 12 hex st hex into dump 18` treats `ES:BX:12` as a pointer to a `FILE` structure, treats the `CLOCK` field within that structure as the far address of one or more bytes, and then hex dumps 18 bytes at that address.

### Log Your Machine's Activity

Out of a `NDOS.MCR` comes a nearly ready-made module application for `INTRSPY`. Using one section of `INTRDOS.MR` as a starting point, you can write a script `EXEC.MCR` that will maintain a log of all the programs run on a computer, together with their command-line arguments, starting time, and completion codes. This script is shown as Listing 5-9.

#### Listing 5-9: EXEC.SCR Logs Program Activity

```

; EXEC.SCR
structure param_block fields
  env_seg (word,hex)
  args (dword,ptr)
Intercept 21h
function 0ah
  on_exit output (ds:dx[1]->byte,string,60)
function 4bh
  on_entry
    output
      (0,046Ch->dword,dec) " " ; time
      (ds:dx->byte,asc)>2,64) ; executable
      (es:bx->param_block args->byte,string,32) ; cmdline args
  on_exit
    if (cflag == 1) sameline " [FAIL - ax]"
function 00h 4ch
  on_entry sameline " [RET - at]"
function 31h
  on_entry sameline " [TSR - at]"
Intercept 20h
  on_entry sameline " [RET20 - at]"
Intercept 27h
  on_entry sameline " [TSR27 - at]"

```

This script reports on all calls to the `DOS EXEC` function, and on all calls to the surprisingly large number of `DOS` functions and interrupts that handle program termination. `INT 21h` function `0Ah` (where `4ch` was originally) is intercepted so that the actual command line typed by a user can also be inspected. In addition, the function `0Ah` can certainly be used by programs that use `COMMAND.COM`. `EXEC.MCR` also displays the `ROM BIOS` time taken by each process, which displays how much time was spent in each program. Here is sample output from `EXEC.MCR` after running `HELLO.C` with Microsoft `C 6.0`:

```

C:\Hello c
534139 C:\MSCBIN\c1.exe hello c
534150 C:\MSCBIN\c1.exe [RET 00]
534194 C:\MSCBIN\c2.exe [RET 00]
534263 C:\MSCBIN\c3.exe [RET 00]
534312 C:\BIN\c.exe wd \c:\temp\0049511k [RET 00] [RET 00]
cmdspy report
536557 C:\UNDOC\MAXEXE\CMDSPY\EXE report [RET 00]

```

to the number of seconds spent in a program, subtract its ticks, and, for example, 534 263 - 534 150 = 118 ticks. If you start the wat program, then divide the result by 18.2 (the number of timer ticks per second).

THE SCR program is a case of a case, with lots of INTRSY. Notice that all the termination functions use SAME LINE so that the program termination status will be printed on the same line as the invocation. However, on a program that spawns a program as CHILD, the status of the child program will not appear on the appropriate line. A future version of INTRSY should amend that, in which case the THE SCR code would have the SAME LINE replaced with IPCHILD. INTRSY did do an output stack so that results could be globally accessed, but THE SCR accomplishes a similar function for which others have no doubt spent days writing complete errors. The THE SCR script took only a few minutes to cook, not a lifetime.

### Monitoring Disk I/O

It's always fun to look someone's desktop, but after you've read it, imagine how to get it would take to write it yourself.

The idea behind DISK SCR is to develop DOS file system calls made in specific processes, as well as to log the related BIOS disk activity generated by such calls. This script uses the BIOS command to run a specific program that you want to monitor. This helps you watch, as an example, FORMAT.COM, without inadvertently also watching COMMAND.COM.

### Listing 5-10: DISK.SCR Watches Disk Related DOS and BIOS Calls

```

* DISK.SCR
* This script relates DOS disk calls to the hard disk
* BIOS calls involved
*
* DOS 4+/Compaq DOS 3.31+ =>32M partition
structure big fields
  sector (dword,hex)
  num (word,hex)
  addr (dword,ptr)
Intercept 21h
function 32h
  on_entry output "2132: Get DPM drive " dt
; function 40h
  * disk M (e.g. just to see messages (e.g. from FORMAT)
  * on entry if (bs = 1) , stdout
  * output (ds dx=>byte,4c*12,c)
function 44h
  subfunction 09h
  on_entry output "214409 IOCTL drive " bt Remote?
  subfunction 0ah
  on_entry output "21440A: IOCTL drive " bl
  if (cl = 40h) sameline ["40 Set Device Parameters"]
  * (cl = 42h) sameline ["42 Format and Verify Track"]
  if (cl = 60h) sameline ["60: Get Device Parameters"]
  subfunction 0fh
  on_entry output "21440F IOCTL Set Logical Drive " bl

```

```

        on_exit if (cflag == 0) sameline " => " at
function 60h
    on_entry output 2160: (anon " (ds:si->byte,asciiz,32) " =>
    on_exit sameline (es:di->byte,asciiz,32)
Intercept 25h
    on_entry
        output "25: Abs Disk Read drv " al ", at sectr "
        if (cx == 0FFFFh)
            sameline (ds:bx->big sector) ", "
                (ds:bx->big num) " sctrs"
        if (cx != 0FFFFh)
            sameline dx ", cx sctrs"
    on_exit if (cflag==1) sameline " [fail]"
Intercept 26h
    on_entry
        output "26: Abs Disk Write drv " al ", at sectr "
        if (cx == 0FFFFh)
            sameline (ds:bx->big sector) ", "
                (ds:bx->big num) " sctrs"
        if (cx != 0FFFFh)
            sameline dx ", cx " sctrs"
    on_exit if (cflag==1) sameline " [fail]"
Intercept 13h
function 0 on_entry output 1300 Recalibrate drive dl
function 1 on_exit output "1301 Disk system status al
function 2
    on_entry
        output "1302: Read " al " sctrs. drv " dl ", head " dh
            ", sectr " cl ", trk " ch
    on_exit if (cflag==1)
        sameline " - FAILED (" ah ")
function 3
    on_entry
        output 1303 Write al " sctrs. drv " dl ", head " dh
            ", sectr " cl ", trk " ch
    on_exit if (cflag==1)
        sameline " - FAILED (" ah ")
function 4
    on_entry
        output "1304 Verify " al " sctrs. drv " dl ", head " dh
            ", sectr " cl ", trk " ch
    on_exit if (cflag==1)
        sameline " - FAILED (" ah ")
function 5
    on_entry
        output "1305: Format " al " sctrs. drv " dl ", head " dh
            ", sectr " cl ", trk " ch
    on_exit if (cflag==1)
        sameline " - FAILED (" ah ")
function 8
    on_entry
        output "1308: Get drive params for " dl
    on_exit
        if (cflag==1) sameline " - FAILED (" ah ")
        if (cflag==0)
            output "Type " bl ", " dl " drvs, max head " dh
                ", max sectr " cl ", max cyls " ch
function Dch
    on_entry output 130C Seek to cyl " ch ", drv " dl ", head " dh
    on_exit if (cflag==1) sameline " - FAILED (" ah ")
function Ddh
    on_entry output "130D Alternate reset drive " dl
    on_exit if (cflag==1) sameline " - FAILED (" ah ")

```

```

function 10h
  on_entry output 1310 Test drive `dl
  on_exit same_line ` - status ` ah
function 15h
  on_entry output "1315: Get type drv = dl
  on_exit
    same_line `:
    if (ah==0) same_line "No disk present : scrs " cx dx
    if (ah=1) same_line floppy Not changed ` scrs ` cx dx
    if (ah=2) same_line floppy changed scrs cx dx
    if (ah=3) same_line "Fixed disk : scrs " cx dx
function 16h
  on_entry output 1316 Get media change drv ` dl `
  on_exit
    if (ah==0) same_line "Unchanged"
    if (ah==6) same_line "Changed"
function 17h
  on_entry
    output "1317, Set type drv ` dl
    if (al=0) same_line no disk`
    if (al=1) same_line "reg disk in reg drv"
    if (al=2) same_line "reg disk in high dens. drv"
    if (al=3) same_line high dens. disk in high dens. drv`
    if (al=4) same_line "720k disk in 720k drv"
    if (al=5) same_line "720k disk in 1.44M drv"
    if (al=6) same_line "1.44M disk in 1.44M drv"
function 18h
  on_entry
    output 1318 Set media type drv dl scrs/trk cl
    , trks " ch " :
  on_exit
    if (ah==0) same_line "0E"
    if (ah==1) same_line "Not available"
    if (ah==0ch) same_line "Not supported"
    if (ah==80h) same_line "No disk in drive"

```

```

run 11 12 13 14
report

```

Note how the script jumps the results to the screen. It is sometimes useful to see the output in hexadecimal and it is also fun to view it in a editor. To capture the results to a file simply run `CMDSPY REPLY=1 FORMAT.CM` after the command completes.

Figure 5-10 shows a small portion of the output from `DISK SCR` that results from formatting a double density 5.25" diskette. You can see that there is also a small amount of output that generated by `FORMAT` itself. The first four lines of the `CMDSPY REPLY=1` also show `DCS` loading `FORMAT.COM` from the hard disk (drive 80):

#### Figure 5-10: INTRSPY DISK SCR Watches FORMAT B:

```

1302 Read 01 scrs drv 80, head 03, scrs 06, trk 71
1302 Read 01 scrs: drv 80, head 03, scrs 10, trk 00
1302 Read 40 scrs drv 80, head 03, scrs 06, trk 71
1302 Read 01 scrs drv 80, head 04, scrs 07, trk 71
21440: 10CF drive 02 [60 Get Device Parameters]
21440F: 10CF: drive 02 Remote?
21440F: 10CF: drive 02 Remote?
2160: Canon B.COM => B:/COM
21440b: 10CF, drive 02 [60 Get Device Parameters]
21440F: 10CF: Set Logical Drive 02 => 00
Insert new diskette for drive
B

```

and press ENTER when ready...

```

21440b: IOCTL drive 02
21440b: IOCTL drive 02
Checking existing disk format.
25: Abs Disk Read drv 01, at sectr 0000, FFFF sctrs
1302: Read 02 sctrs drv 01, head 00, sctr 01, trk 00
Saving UNFORMAT information.
2132: Get 0PB drive 02
.. details omitted .
1316: Get media change drv 01- Unchanged
25: Abs Disk Read drv 01, at sectr
1302: Read 01 sctrs drv 01, head 01, sctr 01, trk 00
25: Abs Disk Read drv 01, at sectr
1302: Read 02 sctrs drv 01, head 00, sctr 08, trk 00
25: Abs Disk Read drv 01, at sectr
... details omitted
21440b: IOCTL drive 02 (40 Set Device Parameters)
21440b: IOCTL drive 02 (42 Format and Verify Track)
1318: Set media type drv 01 sctrs/trk 0F, trks 4F; OK
21440b: IOCTL drive 02
130A: Verify 0F sctrs: drv 01, head 00, sctr 01, trk 00
0
percent completed.
21440b: IOCTL drive 02
130A: Verify 0F sctrs: drv 01, head 01, sctr 01, trk 00
1
percent completed.
21440b: IOCTL drive 02
130A: Verify 0F sctrs: drv 01, head 00, sctr 01, trk 01
.. details omitted .
100
percent completed.
Format complete
26: Abs Disk Write drv 01, at sectr 0000, FFFF sctrs
1303: Write 01 sctrs drv 01, head 00, sctr 01, trk 00
.. details omitted .
21440b: IOCTL drive 02 (40 Set Device Parameters)
2132: Get 0PB drive 02
1302: Read 01 sctrs drv 01, head 00, sctr 01, trk 00
1302: Read 01 sctrs drv 01, head 00, sctr 02, trk 00
1302: Read 01 sctrs drv 01, head 00, sctr 01, trk 00
.. details omitted

```

Whereas DISK SCR groups functions together by interrupt and function number, the output resulting from running the script together with the DOS `FORMAT` command is quite different. Here you see the listing of BIOS calls that DOS calls. In Figure 5-10 you can clearly see how `FORMAT` displays a Fuel Gauge mode "X" odometer, which we're able to show here in most morning IN, 20, F, 000 40b. If the real DOS INT 21h function 44h Subroutine 010h is to format and verify a track, then you indirectly call BIOS INT 13h Functions 18, 05h, and 04h to format and verify 18 sectors that make up the track. After the format itself is complete, as the above INTRSPY output shows, to create the FAT and root directories on the newly formatted disk `FORMAT` calls the DOS Absolute Disk Write interrupt INT 26h. This interrupt in turn indirectly calls to BIOS Write Sector Function INT 13h Function 05h. In brief, DOS calls a disk driver, which in this case calls the BIOS (see Chapter 8).

DISK SCR handles INT 25h and INT 26h calls on systems with partitions larger than 52,000 bytes. This function is important even when you are formatting a floppy disk because even then,

FORMAT uses the alternate form of INT 26h where CX holds the value FFFFh and DS:BX points to a structure that is DISK SCR is called BIL. You can also see from this display that FORMAT uses two undocumented DOS calls: Resolve Path String to Canonical Path String (INT 21h Function 60h) and Get DDP (INT 21h Function 32h). Again having INTRSPY means that you can do this kind of exploring without disassembling the code.

## MEM

One last INTRSPY script worth examining is MEM.MEM Listing 5-11, which in 24 lines of INTRSPY code can monitor all DOS memory allocation by intercepting INT 21h Functions 48h (Allocate 4K Block) and 4Ah (Resize). Again, Function 4Ah (FAIL) is monitored as well, so that you can watch a program performed the memory operation. Chapter 7 presents another more comprehensive version of MEM.

### Listing 5-11 MEM.SCR Watches DOS Memory Management Calls

```
.MEM.SCR
intercept 27h
function 48h
  on_entry
    output "ALLOC " (bx, dec) " paras"
  on_exit
    if (cflag==1)
      sameline " FAIL (" (ax, dec) " ), only " (bx, dec) " available"
    if (cflag==0)
      sameline " - seg " ax "h"
function 49h
  on_entry output "FREE seg " ax "h"
  on_exit if (cflag==1) sameline " denied (" (ax, dec) "h)"
function 4ah
  on_entry
    output REALOC seg ax "h to (bx, dec) paras
  on_exit
    if (cflag==1)
      sameline " FAIL (" (ax, dec) " ), only " (bx, dec) " available"
function 4bh
  on_entry output (ds:ax->byte,asciiz,"04")
```

MEM can be used to track down memory allocation bugs, but it is also useful in a hands-on examination of the DOS memory allocation services discussed in Chapter 3 of this book. For example, you can see immediately that DOS programs typically are allocated all available memory:

```
C:\WINDOWS\SYSTEM32\EXE
ALLOC 65535 paras FAIL (81), only 37705 available
ALLOC 37705 paras - seg 1610h
FREE seg 1610h
```

SPY does not allow me to show all the uses that even my limited imagination has found for INTRSPY.

## Writing a Generic Interrupt Handler

Let's step back now and discuss some of the design issues behind INTRSPY. Some of this sounds like a bit of a speculation for INTRSPY because, in fact, it does come from the functional specification initially drawn up for INTRSPY.

Traditionally, DOS programmers would write a host of small, tailor-made programs to perform the kind of operation on we have been doing. In the past, I have written separate programs to monitor NetBIOS (INT 5ch) calls, DOS (INT 21h) calls, and FMS (INT 6h) calls. The cycle has been to write a simple INR.dbug.1 and eventually have something that could print or write a register val

ues into or out of intercepted functions. In principle, that initial version cost a few weeks of original programming and debugging effort. Subsequent versions that refined the list of functions being monitored to that which added a new or increased capability came easier and more quickly, but they still required recompilation and debugging.

All monitoring programs are essentially the same, however, whether they monitor DOS, NetBIOS, FMS, or something else. This suggests that it would be possible to write a generic monitoring program. By extension, this means writing a generic interrupt handler. The real data intercepts, functions, or subfunctions can wish to monitor, would be parameters to the generic interrupt handler. To ensure that handling the traps doesn't detract or distract from some investigations, the trap handler will accept parameters, without the necessity of debugging, as a high-priority, low-level number of parameters that need to be available readily. However, and the need for a broad ranged capabilities suggest that a commercial, basic, switch-based ISR is unlikely to be adequate. A script interpreting precompiled interrupt handling code would be better.

Because DOS uses commands and different interrupt numbers and provides its services through functions and subfunctions within these interrupts, generally specified via INT and INT\_ and because you would want to be able to monitor these interrupts and their function calls selectively, INTRSPY must be capable of monitoring any subfunction of any function of any interrupt. It should therefore allow you to build some user-originate and exploration applications, not just in the field of undocumented DOS APIs but in a whole range of DOS-related functions DOS related interrupt-based services. Logging DOS activity in a log would be a script, for example. Reading DOS disk access to the underlying device driver is a DOS interrupt, watching FMS, is a NetBIOS session. All these functions should be within reach simply by using different, easily obtainable parameters. INTRSPY should be a platform-independent debug tool for 386 software developers.

The preliminary specification is beginning to form. The program should be script-driven. It is allows for transportable, repeatable, and repeat experiments and debugging tools. The program should intrude as little as possible. It should not make no DOS calls, should not hook to interrupts, and should consume little memory. An interrupt that is generated in a session that is suspended, any interrupts that are intercepted, should not lose that it supposed to be intercepted. These factors constitute noise and can affect some trials in many ways, the systems being studied. If CMIOSPY made undocumented DOS calls, for example, it should show a path to the report, the NDBG, NDR. That it is perhaps ironic that it is work devoted to the aid of someone, an investigator, DOS, as a critical tool in the development of ISRs and system-level software, should not have any role in developing CMIOSPY and INTRSPY, with as few undocumented DOS calls as possible. Once independent, INTRSPY generates no interrupts and uses no DOS services, directly.

The program should provide dumps of register contents, flags, and DOS and other structures in memory. It should be able to register capabilities was specified initially. It should use the DOS list of functions here and the associated DOS function IDs (see INTRINSISER previously). As was noted, however, once an initial capability set of traps has been provided, handling a single, thread, can make it useful in many, originally, interesting ways. As was a read said, the central requirement is to be able to look, from session to session, on different subfunctions of different functions or different interrupts, reconfigure it should be, easy and painless so that it does not impede learning, experimentation, or the debugging progress.

Access should be available both before and after an interrupt is serviced. The program needs to be able to see and set source parameters that the caller supplies, as well as the structures, registers, and flags that the interrupt function returns.

A non-volatile database of pointers should compile the script, handle the off to the reader's portion, and describe output. When we were first discussing the specification for INTRSPY, we had been using some of the other, available, interrupt monitoring software packages. It was not clear that any of them offered the flexibility we needed.

The options here were either a shell implementation or a TSR (transient controller) combination. In a shell implementation, programs under scrutiny would be run from within the system. Much like those of a debugger, a shell implementation's facilities would be available until the user quit the user interface. The problems with this approach include potential for much reduced memory availability for the program being studied and less flexibility to run commands and batch files as well as programs. Besides, it robs control over the environment and the ability to limit monitoring to a specific run of a particular program. Thus, INTRMPY includes this as an option with the RUN statement. We decided to go with the TSR (transient controller) combination. A relatively small TSR would perform the monitoring function, and a separate, less memory-constrained program would perform script compilation, script formatting, and printing; it would handle all communications with the TSR. We saw this approach as preferable because the transient portion itself can be made to act as a shell.

This approach does have its disadvantages, however. Complexity is introduced through the decomposing of the compilation from the execution of the script. This is awkward only because, unlike normal compilation and execution cycles, script source is needed to be able to decode the results of the execution.

The results are posted on-screen rather than write them to disk or pop up screens. This removes the need to plan for file I/O within the generic interrupt service routine (ISR) code. File I/O by itself generates interrupts, changes the state of the operating system by a little or a lot, and leads to coding complexity when implemented properly in ISRs (as is shown in Chapter 9).

### **The Problem with Intel's INT**

Now that we have explained how to use INTRMPY and discussed some of the thinking behind it, we need to discuss, briefly, some problems with the Intel INT instruction. These problems stand in the way of anyone seeking to write a generic interrupt handler. The INT instruction is the source of a great deal of the flexibility of the PC architecture because the ability to get and set interrupt vectors means that system services, including DOS itself, are virtually extensible, replaceable, and monitorable. Yet, paradoxically, the INT instruction is also remarkably inflexible in two key ways:

- A user, trap handler, does not know which interrupt number invoked it.
- The INT instruction itself expects an immediate operand; you cannot write MOV AX, 21h and then INT AX; you must write INT 21h.

The first problem raises the question: How will the program trap a variable number of user-specific interrupts? There are at least three possibilities here:

The first possibility is to use a generic interrupt service routine for all interrupts the user specifies. When invoked by an ISR, it uses the return address on the stack to find out what INT instruction issued the interrupt.

This would be ideal by an elegant, economical solution. Unfortunately, it is an unreliable strategy because many high-level language compilers—and the important Simulate, Int and Exec, Int functions of Windows 1.0—do not make complex interrupts into PUSH and far CALL instruction sequences, since that does not affect INT. Others push the address of the handler on the stack and RET to it. In order to overcome this difficulty, to simulate an INT instruction, a generic ISR would need a small disassembler.

One reason for a failure of either way of performing what should simply be an INT is in fact the second problem—that the INT instruction itself can't be parameterized. Thus, different compiler vendors implement functions such as simint differently.

For other trapping facilities, a more reasonable implementation is that it relies on the caller's stack being large enough for sequential processing. Bear in mind that the program is going to have to be able to handle DOS internal functions, calls for which DOS will have switched to its own small stack of less than 400 bytes (for exact sizes, see the appendix entry for INT 21h Function 5FD06h).



A second possibility involves coding 256 small stubs. Only those interrupts the user specifies would be redirected to the appropriate stubs. When invoked, a stub would record its interrupt number, save away the caller's stack, switch to the program's internal stack, then call the generic ISR. This is a better solution—but still would use several kilobytes of memory.

There are variations on these, but none were appealing, and I decided to go into over a drought was at least a more interesting solution—that is, to allocate a custom ISR object on the heap for each interrupt to be monitored. I selected an object because it was a structure containing machine code to perform stack switching, registers for saving state, and the specific compiled code associated with the interrupt processing. INTRSPY contained a skeleton ISR that was copied onto the heap and filled up with some non-0 machine-independent addresses. The interrupt processing was then actually performed by procedures called from the ISR machine code on the heap.

Will the program base code copy the entire data structure objects at runtime because of allocating on the heap and therefore cannot expect every ISR structure to be paragraph aligned? No. Each ISR starts on a paragraph boundary, wasting up to 15 bytes to ensure it. This amount is insignificant and allows the entry point as well as the data fields in the ISR structure to be at a constant offset. INTRSPY can thus be viewed as a program loader in which the programs are just small pieces of ISR code and the transfer space substitutes for DOS memory.

How would the interrupt processing instructions be stored. My first reaction was to think in terms of compilation instructions originating instead for the first system I implemented a linked tree of small structures that describes the activities in the processor. The resident ISR processor in version 1.0 walked the linked list of function records stored for the interrupt that has been intercepted, processing the subfunction branches of each. Each subfunction record was a linked list of subfunction records, each of which pointed to a "before" and "after" branch, whose branches were in turn, linked lists of records that specified the conditionals and resulting "at" storage to be performed before and after the interrupt is serviced.

### Changed Implementation in INTRSPY 2.0

Version 2.0 was born out of experience with version 1.0, which is as it should be. A number of new early decisions have been overturned, and some additional thoughts have led to other design changes.

For example, I now realize that I should probably have interposed much of machine code generation right from the start. Many of the powerful newer script language features have required less incremental effort to implement than they would have from the database approach in INTRSPY 1.0. Indeed many would simply not have been feasible with that approach.

It also became clear that INTRSPY needed to be a little more than a simple in-memory, no-any-disk-generated-and-stored-on-a-MDISPY. Most of the conversation between CMDSPY and INTRSPY was unnecessary. INTRSPY did not need a hand in it. Instead CMDSPY needed that INTRSPY as a block of memory that it could act on directly. This approach produced a dramatic drop in the size of the case, that is, ignoring another result space allocation. INTRSPY memory footprint from 20K in 1.0 to less than 7K in 2.0.

Note that one of the reasons I did not want to know the undocumented details between INTRSPY and CMDSPY is under development their own CMDSPY replacements. This was documented in Ralf Brown's and Tim Riles's *PC Interrupts* and that we had "undocumented stuff" in *Undocumented DOS*—but now most of this goes away. The interface is 3.0 is simple.

### Implementation

Let us take a look at the structure of INTRSPY and its transient controller CMDSPY. The system functions as follows:

- INTRSPY is loaded and reserves some memory for holding intercept code and results data.



Int 25h and 26h (i.e. Abort to Sector Read/Write interrupts) must also be specialized. Even though they are accessed through a different interrupt mechanism, they are implemented as far as routines. This means that if the flags that were pushed onto the stack by the int 25h instruction are still on the stack and the SI is not as it was when the instruction was issued, therefore the stack must be adjusted by the caller.

Another DOS-specific complication that we saw when reviewing C-MIDSPY messages involves the attempt to restore the EIP (EIP handler). As we saw, C-MIDSPY itself is running as part of a DOS EXEC function. If the user's interrupt handler is in mid interrupt, restoring the vector and performing a replacement of the ISR with another one would be disastrous since the DOS EXEC function would return to an address that had been superseded.

The key theme to remember here is that all this complication is encapsulated within the generic interrupt handler. INTRSPY's success in INTRSPY scripts happens without agonizing over the usual complications of interrupt handling. Success in INTRSPY seems to make these details less a key benefit to using such an interrupt handling language.

Because INTRSPY's become a commercial product, source code is not provided. If you are interested in getting a general idea of how an interrupt monitoring program works, you might want to look at *Body Builders and Methodologies: Collecting Program Statistics with INFORMER*, by PC Magazine, C. Lee, Ann Z. P. Jones Press, 1993.

## The Future of INTRSPY

INTRSPY's success as a standard and a robust and widely used debugging tool. Pressure for additional features and support has been growing, and although 2.0 is a considerable advance over 1.0, a still more sophisticated package is already planned for the first commercial release of INTRSPY version 3.0. A partial list of the new features for a final version that release includes:

- A full-fledged algebraic language, including C, incorporating support for named variables, math, and a powerful set of built-in functions.
- An enhanced user interface, including static script generation tools.
- Secondary monitor capabilities for real-time monitoring of events.
- Additional remote script options, including XMS files and transmission via serial or network connection to a second PC.
- Enhanced report and screen formatting capabilities.
- Full documentation.

For more information, please contact me through my CompuServe address: "0401 315".





Since programmers frequently have questions about the legitimacy of disassembly, this chapter also briefly discusses the law surrounding reverse engineering and trade secrets.

Of course, there is more to DOS than just IO.SYS and MSDOS.SYS. There are also external programs such as COMMAND.COM, MSDOSHELP, and PRINT.COM, which is probably the most heavily disassembled DOS utility and the one on which many TSR writers first figured out their craft.

Whether or not your disassembly DOS depends of course on what interests you. The explanation of the INT 21h dispatch code in this chapter may prove useful to you, even wanted to know about how DOS works internally. On the other hand, if you assemble primarily to just know exactly what is going on inside MS-DOS and you have the money to pay for this information, you may want to license Microsoft's DOS OEM Adaptation Kit, which includes assembly language and C source code for many parts of DOS, as well as CRT files with full symbol information for those parts where direct symbol code is not provided. We take a quick look at the OAK contracts later on.

## What is MS-DOS?

MS-DOS is a file allocation methodology. Everyone knows what it is when they see it, but almost no one can define it.

The real MS-DOS is what the C prompt shows. What that ostensible user interface seems practically irrelevant to MS-DOS is a more certainly necessary part of DOS. The C prompt is provided by COMMAND.COM, which is discussed in Chapter 10; except for its more detail, anyone can cash register. As indicated in the last sentence in CONIO.SYS, COMMAND.COM is just a shell around the DOS kernel. Other shells such as 4DOS or the MS-DOS shell are widely available. Get rid of COMMAND.COM, and you still have MS-DOS.

From a programmer's perspective, MS-DOS seems like a collection of INT 21h functions that has been organized into a shell. While the INT 21h functions are the most important service provided by DOS, Windows NT 2.11 may not seem so important. Several application wrappers in Chapter 2 (strings 2.20 and 2.21) and Chapter 10 show how easy it is for a normal program to huddle with INT 21h calls before an application DOS and gets a more elaborate piece of code handles INT 21h doesn't necessarily make it part of DOS.

So 4DOS or whatever is the C prompt or the INT 21h interface, what then is it? And where is it?

The most good is difficult to answer, except to note that DOS is in many ways what text books on operating systems call a microkernel. DOS provides a small bare minimum of services on top of which other more sophisticated services can be built. Much of DOS as a software motherboard into Intel's case is free to programers and users. These extensions come not only from Microsoft but also from key third-party vendors such as Novell, Quarterdeck, Quantas Systems, Centra-Point, and Pico. For DOS—the actual kernel shell of these companies' products must both compete and work together.

Well, that was vague enough!

Microsoft's answer, set at C:\MS-DOS, is easy to answer. MS-DOS consists of two files, IO.SYS and MSDOS.SYS. In both IBM PC-DOS and Novell's DR-DOS, these files are called IBMBIOS.COM and IO.BIOS.COM. Despite the SYS file names, these are not device drivers but binary images. In MS-DOS 6.0, there is a third file, MSDOSHELP.BIN, which Microsoft generally considers a full-fledged third member of the DOS kernel. The SYS and IO.BIOS commands in DOS 6.0 copy MSDOSHELP.BIN over to a floppy along with IO.SYS and MSDOS.SYS. Like these two or three files, a file system DOS, of course, such as FAT and a shell such as COMMAND.COM in order to get much work done.

A very subtle change, MSDOS.SYS controls the DOS dispatch function, which is DOS's handler for INT 21h calls. There are other DOS functions such as INT 25h, 26h, and 2Fh, that MSDOS.SYS and IO.SYS handle as well.

IO.SYS consists of two parts: a loader (MSLOAD.COM) and BIOS support code (MSBIO.BIN). MSLOAD.COM creates IO.SYS by concatenating these two files:

```
copy /b msload.com+msbio.bin io.sys
```

IO.SYS is *not* the BIOS, as looking at DOS programming books might claim, but merely the DOS *interface* to the BIOS. IO.SYS controls the standard device drivers such as CON, AUX, LPT, and COM (see Chapter 7). These device drivers are implemented using DOS calls. For example, the CON driver built into IO.SYS (not precisely MSBIO.BIN) makes INT 10h and INT 16h calls to the ROM BIOS video and keyboard routines.

The MSLOAD.COM part of IO.SYS contains a famous set of routines called SYSENTRY, which is responsible for the bootstrapping of DOS.

We won't discuss SYSENTRY here, as it has already been covered elsewhere (see "How MS-DOS Is Loaded" in Chapter 2 of *Re-Discovering the Secrets of MS-DOS Programming* and "The Components of MS-DOS" in *Discovering MS-DOS for Dummies*). And practically every other book on DOS programming seems to skip past this same basic information on SYSENTRY. Presumably this is not just because the bootstrapping of DOS is an interesting subject, but also because Microsoft already dominates SYSENTRY with DOS.DAK. Chapter 4 provides a thorough, original and useful description of DOS startup in *MS-DOS Internals*, Chapter 1, "The System Configuration"; 2, "The System Initialization"; and 3, "The Startup Sequence". For example, Chapter 1 is first in order to make the subject of SYSENTRY and the list of lists structure (those actual names in the DOS source code) is in *IO.SYS*.

So the DOS boot sequence is not well known. What has not been provided before (anywhere) is any description of how DOS looks like once it's up and running. This primarily requires a description of DOS's INT 21h interface and the INT 21h dispatch table. In other words, for someone who can make an INT 21h call to DOS, sources of DOS programming books or manuals describe what they do, but DOS internals books or the descriptions here are only a glimpse of how an internal user would know how to use the rest of DOS's stack structure. *Microsoft's MS-DOS For Developers* (pp. 352-35) describes the DOS function call mechanism itself. This seems far more important than providing an accurate standard description of how DOS boots up or how SYSENTRY moves segments around in memory.

One of our early exercises was to look at the boot sequence by NOCOW.EXE (see "How MS-DOS 6.0 and up" in the next section) and IO.SYS uses the pre- and DIBSPACE.BIN. And, as DOS 7.0 (Chicago) or COMMAND.COM contains the setting DOS\_FILENAMEID, there is code in IO.SYS that reads DOS386.EXE, which is a big executable, similar to WIN386.EXE.

### Disassembling IO.SYS and MSDOS.SYS

The choice between disassembling SYSENTRY or describing the INT 21h interface is important not because the problem of DOS booting only is interested in looking at largely abstracted low-level goals about booting DOS.

I look at DOS booting only to show that there is a way to acquire the DOS.DAK, which provides assembly language source code to IO.SYS, including the SYSENTRY modules, and you have to disassemble the actual IO.SYS and MSDOS.SYS files on disk. These files are hidden system files, which, however, can be easily unhidden.

```
C:\WINDOWS\CHAP6>attrib -h -s *.sys
```

IO.SYS is about 32k, and MSDOS.SYS is about 37k. Once unhidden, these two files can be disassembled (examined) or reassembled (recreated) in the previous DEBUG (either 0.99 or comes with DOS). After running ATTRIB to unhide MSDOS.SYS or IO.SYS (perhaps not needed), their size, DEBUG loads the file at address 100h, so add 100h to the file size (convert it to a code value) to yield the disassembly address. For example, if MSDOS.SYS is 310669282h bytes,

```
C:\MSDOS2\CHAP6>type msdos.scr
```

```
0100 9382
```

```
B
```

```
C:\MSDOS2\CHAP6>debug \msdos.sys < msdos.scr > msdos.lst
```

The result of `MSDOS.LST` is about one megabyte in size if you use a disassembler such as `SOFT-ICE` to look at the `MSDOS`. In some ways, the output from such a straightforward disassembly of `MSDOS.SYS` looks quite useful. For example, you can quite plainly see `DOS's` `INT 21h` handler speaking to the `IOPL` function number in `BIOS`. This is the `DOS` code called whenever a program generates an `INT 21h`.

```
0A76 0408 FA          CLI
0A76 040C 80FC61       CRP AH,6C          ; is function > 6Ch?
0A76 040F 77D2       JA 03E3           ; yes: error
0A76 0411 80FC33       CRP AH,33
0A76 0414 721B       JB 042F
0A76 0416 74A2       JZ 03BA
0A76 0418 80FC64       CRP AH,64
; etc
```

Likewise, the `MSDOS.SYS` `INT 21h` handler is a reasonable `IOPLSYS` has its own `INT 21h` handler, and if you list the `IOPLSYS` code fragment below, you can see the `INT 21h` handler in `MSDOS.SYS` in the context of `IOPLSYS` using a hard-wired address.

```
1C55 0789 FB          STI
1C55 078A 80FC19       CRP AH,19
1C55 078D 750A       JNE 07C9
... Go to 078Fh if an INT 21h call belonging to an external
... program such as a redirector, SHARE, or NLSFUNC, ends up
... in MS-DOS. This means the external program isn't loaded
1C55 078F 0A7D       DF A,,A          ; is A<0?
; error handling
1C55 07C9 80FC10       CRP AH,10          ; INT 21h AH=10h? (SHARE)
1C55 07CC 74F1       JT 07BF          ; got here, so SHARE not loaded
1C55 07CE 80FC14       CRP AH,14          ; INT 21h AH=14h? (NLSFUNC)
1C55 07D1 74EC       JT 07DF          ; got here, so NLSFUNC not loaded
1C55 07D3 80FC12       CRP AH,12          ; INT 21h AH=12h?
1C55 07D6 7503       JNE 07D8
1C55 07D8 199701       JMP 0972          ; handle DOS internal functions
1C55 07DB 80FC16       CRP AH,16          ; INT 21h AH=16h? (Windows)
1C55 07DE 74D0       JT 07E8          ; might be Windows broadcast
1C55 07E0 80FC46       CRP AH,46          ; INT 21h AH=46h?
1C55 07E3 7503       JNE 07E8
1C55 07E5 193E01       JMP 0926
1C55 07E8 EA05007000      JMP 0070 0005    ; see if IO SYS wants it
```

If you look at the `IOPLSYS` code, after a few minutes it becomes clear that the quality of the disassembly is considerably better than `MS-DOS`. Much better versions of these `INT 21h` and `INT 21h` handlers are shown with `SOFT-ICE` and `SOFT-ICE`. For example, the most important part of the `INT 21h` handler uses the `IOPLSYS` function number as an index into a dispatch table:

```
;; previously moved AH func number into BX
0A76 04FE 809FA73E     MOV BX,IBX+3EA73
0A76 0502 36871EEA05    XCHG BX,SS [05EA]
0A76 0507 36BE1EFC05    MOV DS,SS [05EC]
0A76 050C 36FF16FA05    CALL SS:[05EA]
```

Unfortunately, if you now go and look at `MS-DOS`, presumably the address of the all-important `INT 21h` handler in the dispatch table there turns out instead to be perfectly valid looking code at that address, and it is a rather `IOPLSYS` and `DSE` are in this context totally bogus. This isn't a



problem with DBLSPACE.SYS. A straight disassembly on disk of MSDOS.SYS or IO.SYS, even with a more sophisticated assembler, such as SourceJerk, doesn't produce the results you need.

The problem is that the SYSDISK process is described in the MSDOS File of a moving segment around a cylinder, rather than by segment functions. Address ranges of cylinders that won't match up properly is a state described by DOS's disk layout table. In fact, only the core DOS boot-up functions are made easier to disassemble DOS in general, after the DOS boot-up initialization segment is done, which might include the DOS INIT move to the DOS kernel to the high memory area, or HMA, is complete.

The only problem with disassembling DOS out of memory rather than from a file is on disk, but this moves the SYSDISK code, which is discarded from memory when the boot-up is complete. However, as noted earlier, SYSDISK and the DOS bootstrap process are already well adequately covered elsewhere.

Again, you can't expect to see ANY more, forgetting all the practical stuff that DOS does starting in DOS 6.0. Also, you might find SYSDISK to be a little bit of a link problem with the code that uses it, especially if it is by some fixed amount, which is easy to get wrong, or you don't have a code data path. Just look at the code in IO.SYS that preloads DBLSPACE.BIN. If you do, it seems we ought to take a look at this.

## Examining How IO.SYS Preloads DBLSPACE.BIN

It turns out that static disassembly of IO.SYS is actually pretty easy, even though at first glance the results produced by a disassembler such as SourceJerk inadequate. It's true that references to data don't match up with the actual locations of the data in the file, but once you match up a sizeable piece of data in the file with code that references it, you can figure out everything else.

For example, a SourceJerk disassembly of IO.SYS from MS-DOS 6.0 contains the following data item:

```
54BF:8138 5C 44 42 4C 53 50 43 45 db  '\DBLSPACE.BIN'
```

This is followed shortly by code that, based on the surrounding context (the code calls the INT 21h AX=4B03h Load Overlay function), is probably loading DBLSPACE.BIN. However, the code does not reference offset 8138h. Instead, it references CS:3B62h:

```
54BF:8153 0E          push  cs
54BF:8154 1F          pop   ds
54BF:8155 BE 3B62     mov  si,3B62h
```

If you subtract 3B62h from 8138h, you get 45D6h. If the code at 54BF:8155 really is referencing the '\DBLSPACE.BIN' string at offset 8138h, then 45D6h is the amount which you must add to other data references in this version of IO.SYS in order to locate the data itself. To see if this amount is accurate, just look for another data reference, and see if adding the amount onto it yields a likely-looking address. For example, a little further on in the file, IO.SYS produces an error message:

```
54BF:81E9 0E          push  cs
54BF:81EA 1F          pop   ds
54BF:81EB 0A 5823     mov  dx,5823h
54BF:81EE 04 09     mov  ah,9
54BF:81FD CB 25     int  21h ; DOS Services ah=function 09h
; display char string at ds:dx
```

From the helpful comment supplied by Sourcer on how INT 21h AH=9 works, it is clear that 5823h must be the offset within CS of a string. Adding 45D6h to 5823h yields 9DF9h and there, indeed, is an error message:

```
54BF 9DF9 57 72 6F 6E 67 20 db 'Wrong DBLSPACE.BIN version', 0Bh
```

Thus, we really can pick apart IO SYS on disk. This lets us examine the DOS boot process in particular the recent additions such as the preloading of DBLSPACE.BIN in DOS 6 and the apparent ability to pre-load DOS 386 EXE in DOS 7. "Preloading" means that IO SYS loads for and loads these external programs before processing any DEVICE= statements in CONFIG.SYS. Chapter 4 discussed how Stacker 3.1 uses this interface to get itself preloaded under DOS 6. By examining IO SYS, you can see how the interface works.

For example, after calling INT 21h AX=4B03h to load DBLSPACE.BIN, IO SYS looks for a function pointer at offset 14h in DBLSPACE.BIN:

```
54BF 819F E8 FDD6 call LOAD_OVERLAY, subr. does 21/4B03
54BF 81C6 2E C7 06 0387 0014 mov word ptr cs:[387h],14h ; get func ptr from
54BF 81C6 2E 8C 06 0389 mov word ptr cs:[389h],es ; offset 14h
; ; in DBLSPACE.BIN
```

IO SYS saves away the function pointer provided by DBLSPACE.BIN, and then calls it:

```
54BF 81DA DE push cs ; IO SYS passes DBLSPACE.BIN
54BF 81DB 07 pop es ; a pointer to a buffer;
54BF 81DC 8B 056A mov bx,56Ah ; 56Ah+450Ah=4940h (see below)
54BF 81DF 8B 0006 mov ax,6 ; DOS version
54BF 81E2 2E FF 1E 0387 call dword ptr cs:[387h] ; call DBLSPACE.BIN
; ; function ptr
```

```
54BF 8228 8B 0004 mov bx,4 ; subfunction 4
54BF 822B 2E FF 1E 0387 call dword ptr cs:[387h]
```

```
54BF 4940 8 00 db 18h, 00h ; a communications buffer
```

IO SYS also checks for a 2E2Ch signature at offset 12 in DBLSPACE.BIN. A hex dump of DBLSPACE.BIN reveals the presence of this signature:

```
C:\UNDOC? CHAPedump -dos dblspace bin -bytes 32
0000 FF FF FF FF 42 48 41 08 BB 08 01 44 42 4C 53 50 | . .BHA .. DBLSP
0010 9010 41 43 2C 2E E9 B2 59 00 00 EA 41 08 00 00 EA 8B | AC... Y .A...
```

Further discussion of this interface and its possible role in the ongoing battle between Microsoft and Star Electronics appears in Chapter 3. Here, the point is simply that all existing descriptions of the DOS boot process will need to be rewritten to take into account new additions to DOS such as DBLSPACE.BIN (and, in DOS 7, DOS 386 EXE).

Finally, we note that a few other items are needed to execute INT 21h dispatch code, which is executed by the BIOS. It makes DOS call except for other program that hooks INT 21h has control. It intercepts the call to the interrupt. As we see, there are many important aspects to the INT 21h dispatch code: setting stack switching case of the system (PNP incrementing and decreasing the stack), DOS flag handling of critical sections (Ctrl-Break and critical errors), checking for DOS 3.20 or later (DOS 3.20), and special saving for Windows Enhanced mode.

## Interrupt Vectors and Chaining

Tracking DOS internals requires finding the code or PCBs that handles software interrupts such as INT 21h and INT 2Fh. As we just saw, trying to do this with IO.SYS and MSDOS.SYS on disk can produce inadequate results. In memory, however, it seems like it should be easier to find DOS's INT 21h and INT 2Fh handlers. As every PC programmer knows, there is a documented DOS function, INT 21h AH=35h, which returns to ESI:EAX a far pointer to the code that handles the interrupt given in AL.

Finding the current addresses for INT 21h and INT 2Fh is thus a simple matter of calling INT 21h AX=3521h and AX=352Fh and looking at the returned far pointer or vector as a scalar. This can be wrapped up in a simple program to print out interrupt vectors. Add a few extra instructions, such as trying to figure out the owner of each interrupt vector and disassembling some (perhaps) counter-intuitive instructions at the beginning of the interrupt handler, and the result is INTVECT.C, shown in Listing 6-1. Listing 6-2 shows MAP.C, which attempts to figure out owners.

### Listing 6-1: INTVECT.C

```
/*
INTVECT.C
bcc intvect.c map.c
*/

#include <stdlib.h>
#include <stdio.h>
#include <dos.h>

typedef unsigned char BYTE;
typedef unsigned short WORD;
typedef unsigned long DWORD;

#define MK_IN(fp) (((DWORD) FP_SEG(fp) << 4) + FP_OFF(fp))

extern char *find_owner(DWORD in_addr), // in map.c

#define ARPL 0x03
#define IRET 0xCF
#define JMPF 0xEA
#define JNPF 0xEB
#define JMP16 0xE9

BYTE far *get_vect(int intno) // get INT 21h AH=35h
{
    __asm push es
    __asm mov al, byte ptr intno
    __asm mov ah, 35h
    __asm int 21h
    __asm mov dx, es
    __asm mov ax, bx
    __asm pop es
    // return value in BX:AX
}

void print_vect(int intno)
{
    char *s;
    BYTE far *fp = get_vect(intno);
    printf("INT %02Xh %Fp ", intno, fp);
    if (fp == 0)
    {
        printf("unused\n");
        return;
    }
}
```

```

)
e = find_owner(RK_Like(fp));
printf(" 08s %s\n", s);

switch (*fp, // see if first instruction of interrupt handler
      // is anything really obvious
      case ARPL: printf("arpl -- Windows V86 breakpoint"), break;
      case IRET: printf("iret -- MDP function"), break;
      case JNB: printf("jmp %fp",
                    ((BYTE far *) fp) + fp[1] + 2), break;
      case JMPB: printf("jmp %fp",
                    ((BYTE far *) fp) + *(WORD far *) &fp[1] + 3), break;
      case JMPF: printf("jmp %fp",
                    *((void far *) far *) &fp[1]); break;
)
printf(" r\n");
}

main(int argc, char *argv[])
{
    char *end,
    int intno, i,
    if argc < 2)
        for (intno = 0; intno < 256; intno++)
            print_vect(intno);
    else for (i = 1; i < argc; i++)
        print_vect(strlen(argv[i], &end, 16));
    return 0;
}

```

For example:

```

C:\UNDOC2\CHAP6>intvect 21 28 2f 24
INT 21h  C0B0 0942
INT 28h  18D4 0615  PRINT
INT 29h  007D 0762  IO
INT 2fh  1482 D00D  NLSFUNC

```

## INTVECT and Windows

If you run INTVECT without command line parameters, it dumps out the vectors for all 256 interrupts. This is useful, for example, it determines which interrupts Windows Enhanced Mode takes over; you can run INTVECT - TMP TMP, start Windows, run INTVECT - TMP 2 from inside a DOS box, and then use diff or a similar utility to compare the files TMP TMP and TMP 2. The difference between these two files reveals the interrupts that Windows Enhanced Mode hooks using the low memory interrupt vector table (it also hooks some interrupt vectors in the protected mode interrupt descriptor table). Where - points to the previous DOS output, it from INTVECT, and + points to the output under Windows, part of the output from diff might look like this (the complete output also shows changes to INT 0, 3, 8, 10h, 15h, 1C h, 22h, 23h, 24h, 67h, and 68h):

```

< INT 28h  07F9:15AE  SMARTDRV
> INT 28h  FEF8 0862  DBLSSYS$ arpl -- Windows V86 breakpoint

< INT 2fh  1305 0285  DOSKEY
> INT 2fh  1627:02A7  win

< INT 48h  F00D EA97  DBLSSYS$
> INT 48h  FEE1 0C02  DBLSSYS$ arpl -- Windows V86 breakpoint

```

INT 28h is the DOS idle interrupt, and the Virtual DMA Services (VDS) use INT 4Bh. As you can see, INTVECT examines the first byte of an interrupt handler, looking for code such as the ARPL instruction, which Windows Enhanced mode uses as a V86 breakpoint to force a transition from user (Ring 3) code to VMM (Ring 0) code. The seeming location of the Windows V86 breakpoints inside DBLSSYS5 (DoubleSpace) is misleading. This has to do with the way Windows implements V86 breakpoints (see Chapter DOS Internals, Chapter 2).

To build INTVECT, INTVECT.C should be linked with MAP.C (Listing 6-2). MAP.C attempts to provide the only source for each interrupt vector, using code that is explained in detail in Chapter 7. We will use MAP.C with the program, etc. in this chapter, INTCHAIN.C (Listing 6-5). MAP.C is also compiled with DINTSVC to produce a standard output. For example, running MAP.C on my machine happened to produce the following output, which shows that this machine is running DoubleSpace, MSCDEX, SMARTDRV, loaded high, DOSKEY, also loaded high, and XMS and EMM servers.

```
C:\MSDC2\CHAP6>map
0000700 00009A0 IO
00009A0 0001E80 DOS
0001E80 0002010 B:
0002010 0005780 MSBIOUSE
0005780 0007E40 MSCDOS1
0007E40 0012FA0 DBLSSYS5
0012FA0 00131F0 SETVERXX
00131F0 0013670 XMSXXXX0
0013670 0014950 EMMXXXX0
0014950 0018BA0 MSCDFX
0018BA0 002A7E0 MAP
002A7E0 002CBB0 COMMANB
002CBB0 00D2C60 SMARTDRV
00D2C60 00DDBA2 H
00DDBA2 00DE470 J
00DE470 00DF4A0 DOSKEY
00DF4A0 0100000 HMA
```

### Listing 6-2: MAP.C

```
/*
MAP.C
bcc intvect.c map.c
bcc intchain.c map.c
bcc -DTESTING map.c
*/

#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <dos.h>

typedef unsigned char BYTE;
typedef unsigned short WORD;
typedef unsigned long DWORD;
typedef void far *FP;

#ifndef MK_FP
#define MK_FP(s,o) (((DWORD) s) << 16) + (o))
#endif

#pragma pack(1)
```

```

typedef struct {
    DWORD start, end;
    char name[9],
} BLOCK;

static BLOCK *map;
static int num_block = 0;

int cmp_func(const void *b1, const void *b2)
{
    if (((BLOCK *) b1)->start < ((BLOCK *) b2)->start) return -1;
    else if (((BLOCK *) b1)->start > ((BLOCK *) b2)->start) return 1;
    else return 0;
}

typedef struct {
    BYTE type; /* 'M'=in chain; 'Z'=at end */
    WORD owner; /* PSP of the owner */
    WORD s ze, /* in 16-byte paragraphs */
    BYTE unused[3];
    BYTE name[8], /* in DOS 4+ */
} MCB,

#define IS_PSP(mcb) ((FP_SEG(mcb) + 1 == (mcb)->owner)

WORD get_first_mcb(void)
{
    _asm mov ah, 52h
    _asm int 21h
    _asm mov ax, es:[ebx-2]
    // retval in AX
}

typedef struct DEV {
    struct DEV far *next,
    WORD attr, strategy, intr;
    union {
        BYTE name[8], blk_cnt;
    } u;
} DEV,

#define IS_CHAR DEV(dev) ((dev)->attr & (1 << 15))

DEV far *get_nul_dev(void)
{
    _asm mov ah, 52h
    _asm int 21h
    _asm mov dx, es
    _asm lea ax, [ebx+22h]
    // retval in DX:AX
}

int get_num_block_dev(DEV far *dev)
{
    // can't rely on # block devices at SysVars[20h]^
    // walk once through dev chain just to count # blk devs
    int num_blk = 0;
    do {
        if (! IS_CHAR_DEV(dev))
            num_blk += dev-u.blk_cnt,
            dev = dev-next,
    } while(FP_OFF(dev-next) != (WORD) -1);
    return num_blk;
}

```

```

WORD get_umb_link(void)
{
    _asm mov ax, 5802h
    _asm int 27h
    _asm xor ah, ah
    // return value in AX
}

WORD set_umb_link(WORD flag)
{
    _asm mov ax, 5803h
    _asm mov bx, flag
    _asm int 27h
    _asm jc error
    _asm xor ax, ax
error:
    // return 0 or error code in AX
}

WORD get_dos_ds(void)
{
    _asm push ds
    _asm mov ax, 1203h
    _asm int 27h
    _asm mov ax, ds
    _asm pop ds
    // retval in AX
}

/* find IO.SYS segment with built-in drivers */
WORD get_io_seg(DEV far *dev)
{
    WORD io_seg = 0;
    do {
        if (IS_CHAR_DEV(dev))
            if (_strncmp(dev->u.name, "COM", 8) == 0)
                io_seg = FP_SEG(dev); // we'll take the last one
        dev = dev->next;
    } while(FP_OFF(dev->next, 1) != (DWORD) 1);
    return io_seg;
}

static int did_init = 0;

void do_init(void)
{
    MCB far *mcb;
    DEV far *dev = get_nul_dev();
    WORD dos_ds, io_seg, mcb_seg, next_seg, save_link;
    BLOCK *block;
    int blk, i;

    map = (BLOCK *) calloc(100, sizeof(BLOCK));
    block = map;

    io_seg = get_io_seg(dev);
    block->start = io_seg << 4, block->end = (DWORD) -1;
    strcpy(block->name, "IO"); block++;

    dos_ds = get_dos_ds();
    block->start = dos_ds << 4, block->end = (DWORD) -1;
    strcpy(block->name, "DOS"); block++;

    // should really check if there IS an MBR!
    block->start = 0x100000L, block->end = 0x10FFEEh,

```

```

strcpy(block->name, "HMA"); block++;
num_block = 3;

/* walk NCB chain, looking for PSPs, interrupt owners */
if (_osmajor >= 4)
{
    mcb_seg = get_first_mcb();
    mcb = (NCB far *) MK_FP(mcb_seg, 0);
    if (_osmajor >= 5) // be lazy, see ch 7 for DOS < 5
    {
        save_umb = get_umb_link();
        set_umb_link(1); // access UMBA too
    }

    for (;;)
    {
        next_seg = mcb_seg + mcb->size + 1;
        if (IS_PSP(mcb))
        {
            block->start = ((DWORD) mcb_seg) << 4;
            block->end = ((DWORD) next_seg) << 4;
            strcpy(block->name, mcb->name, 8);
            block->name[8] = '\0';
            block++, num_block++;
        }
        mcb_seg = next_seg;
        if (mcb->type == 'M')
            mcb = (NCB far *) MK_FP(next_seg, 0);
        else
            break;
    }
}

/* walk dev ce chain looking for non-builtin drivers */
blk = &dev_num_block dev(dev);
do {
    NCB far *dev_mcb;
    if ((FP_SEG(dev) & dos_ds) && (FP_SEG(dev) != io_seg))
    {
        block->start = (((DWORD) FP_SEG(dev)) << 4) + FP_OFF(dev);
        dev_mcb = (NCB far *) MK_FP(FP_SEG(dev)-1, 0);
        if (dev_mcb->owner == 8)
        {
            dev = dev->next;
            continue;
        }
        if (dev_mcb->type == 'M')
            block->end = block->start + ((DWORD) dev_mcb->size << 4);
        else
            block->end = (DWORD) -1;
        if (IS_CHAR_DEV(dev))
        {
            strcpy(block->name, dev->u.name, 8);
            block->name[8] = '\0';
        }
        else
        {
            blk = dev->u.blk_cnt; // block drivers in reverse order
            block->name[0] = blk + 'A';
            block->name[1] = ':';
            block->name[2] = '\0';
        }
        block++; num_block++;
    }
} while (dev = dev->next);

```



```

} while((FP_OFF(dev >next) != (DWORD) 1);

if (_osmajor >= 5)
    set_umb_link(save_link);

qsort(map, num_block, sizeof(BLOCK), cmp_func);

for (i=0, block=map, i<num_block-1, i++, block++)
    if (block->end == (DWORD) -1)
        block->end = map[i+1].start;
if (block->end == (DWORD) -1) // last one
    block->end = 0xFFFFFFFF;

did_init = 1;
}

char *find_owner(DWORD lin_addr)
{
    BLOCK *block;
    int i;

    if (! did_init) do_init();

    for (i=0, block=map, i<num_block; i++, block++)
        if ((lin_addr >= block->start) &&
            (lin_addr <= block->end))
            return block->name;

    /* still here */
    return (char *) 0;
}

#ifdef TESTING
main()
{
    BLOCK *block;
    int i;
    do_init();
    for (i=0, block=map, i<num_block, i++, block++)
        printf("%08X %08X %s\n",
            block->start, block->end, block->name);
}
#endif

```

With the exception of unused interrupt vectors and those (such as INT 4FH) that point to data rather than code, you can track addresses displayed by INT 13C1 and disassemble them to see how a given interrupt is used. As an example, Figure 6-1 shows INT 29H, which is the end-of-medium flag (EOM) for floppy drives, as located by default in the CHS driver provided by DOS 5.0.

**Figure 6-1: Default implementation of INT 29H**

```

C:\UNB0C2\CHAP6>intvect 29
INT 29h 0070 0762 10

C:\UNB0C2\CHAP6>debug
-u 70 762
0070:0762 50          PUSH  AX
0070:0763 56          PUSH  SI
0070:0764 57          PUSH  DI
0070:0765 55          PUSH  BP
0070:0766 53          PUSH  BX
0070:0767 840E       MOV   AH,0E
0070:0769 8B0700     MOV   BX,0007
0070:076c CB1D       INT  1D

```

```

0070 076E 50          POP     BX
0070 076F 50          POP     BP
0070 0770 5F          POP     DI
0070 0771 5E          POP     SI
0070 0772 5B          POP     AX
0070 0773 CF          IRET

```

This is a straightforward INT 20h here is just a wrapper around INT 10h AH=0Fh, which is the BIOS interrupt to write a character on teletype mode.

Of course, there's a `_____` everywhere that simple. For example, if you install ANSYS, which is a replacement CON driver, INT 29h points somewhere else.

```

C:\WINDOWS\CHAP6>intsect 29
INT 29h  007D:0762

```

```

C:\WINDOWS\CHAP6>undoc2\chap6\devlod \dos\ansi.sys

```

```

C:\WINDOWS\CHAP6>intsect 29
INT 29h  6EB3:0510  DEVLOD

```

Because we installed ANSYS using DEVLOD, the INTVIEW program shows DEVLOD as the source of the interrupt vector; the source of course, is actually the new CON driver in ANSYS. Now the code at 6EB3:0510 is no longer just a wrapper around an INT 10h call. Instead, it directly manipulates the contents of a segment 10000h and contains special handling for ANS escape control codes. So among the code here would take us to far ahead even for a chapter such as this that rambles non-stop and seldom through the BIOS code. The point here is merely that the INTVIEW program simply doesn't help us point DBI to such segment offset addresses to assemble.

But there's a major problem here. Recall that we are interested in looking at the DOS INT 21h and INT 21h handlers. INTVIEW can, of course, print out the addresses of the INT 21h and INT 21h handlers.

```

C:\WINDOWS\CHAP6>intsect 21 29
INT 21h  0F95:3280  MSCDEX
INT 29h  1305:0285  DOSKEY

```

However, as INTVIEW indicates, these interrupt vectors point not to DOS but to DOS add-ins such as MSCDEX and DOSKEY. In fact, it is practically *guaranteed* that except on the latest, freshly installed specific system, which use ALI, USER, BAI, or CONSOLE.SYS file, INT 21h, INT 29h, and many of the DOS interrupt vectors now point into DOS. The INT 21h and INT 21h vectors are pointing outside of the program's binaries rather than to the DOS products and.

Of course, if you're interested to examine MSCDEX's INT 21h handler or DOSKEY's INT 21h handler, the INTVIEW results are sufficient. They provide all the information needed by a debugger such as DBI or SYMDEB, a handy debugger that Microsoft once included with the Windows SDK. For example, by using DBI or SYMDEB to assemble the 1305:0285 address displayed by INTVIEW for INT 21h, we can see that DOSKEY watches for the Windows and task switcher initialization broadcasts, INT 21h AX=1605h and AX=4D05h. DOSKEY clearly uses the same piece of code at offset 0299h to handle both calls. We can also see confirmation that, as documented in Microsoft's *MS-DOS Programmer's Reference*, DOSKEY responds to INT 21h AH=48h calls.

```

C:\WINDOWS\CHAP6>intsect 21
INT 21h  1305:0285  DOSKEY
C:\WINDOWS\CHAP6>debug
~> 1305:0285
1305:0285 300516      CMP     AX,1605
1305:0288 740F          JZ     0299
1305:028A 300540      CMP     AX,4D05
1305:028D 740A          JZ     0299
1305:028F 80FC48      CMP     AH,48

```

```

1305:0292 741B      JZ      02AF
1305:0294 2EFF2E5F02 JMP     FAR CS:[025F]
* ...

```

But if for example we want to see `MSDOS\SYSTEM\INT21h` handler rather than `IO$KEYSYS`, and if `IO$KEYSYS` is loaded after `MSDOS\SYSTEM\INT21h` is of no use. Note however that unlike `MSDOS\SYSTEM` and `IO$SYS` programs such as `MSDOS\EXEC` and `IO$KEYSYS` are easy to disassemble on disk with a program such as `SourceFromAC` communications.

More importantly, `INT21h` doesn't help us get the address of what we might call the "Original `INT21h` Handler" inside `MSDOS\SYSTEM`. Nor does it help with finding the original `INT21h` handlers inside `MSDOS\SYSTEM` and `IO$SYS`.

Why? Because, as we've indicated at a kind of last-in-first-out (LIFO) stack. The point was made at the beginning of this chapter that the DOS philosophy is to provide the bare minimum operating system services along with facilities for "undoing" DOS. As discussed (or greater detail in Chapter 9 on FSR), one of the keys to undoing DOS is `INT21h` API 25h (i.e. DOS Sector Interrupt Vector function). Along with the Get Interrupt Vector function (API 35h) (i.e. Sector Vector function) allows the creation of what are called interrupt chains which are essentially linked lists of LIFO stacks of code. An interrupt chain consists of two or more pieces of code that handle the same interrupt. The following code fragment (adapted from the `FUNCO032` and `IO$SERV` programs from Listings 2-20 and 2-21, illustrates this:

```

void interrupt far *prev(); // ptr to previous handler in chain
prev = _dos_getvect(0x21); // call 21/35 - get previous
_dos_getvect(0x21, my_int21_handler); // call 21/25 -- set new
// ...
void interrupt far my_int21_handler(REG_PARAMS r)
{
    // look at AH to see if we're interested
    // ...
    _chain_intr(prev); // pass interrupt down to previous owner in chain
}

```

The chain method works as far as `INT21h`. The `prev` is interrupt handler in the chain will not return. It is important to note that sometimes `exec` (i.e. `CALL` rather than `JMP`) the previous handler. This is as a handler to post process the interrupt after the previous handler has done its work rather than to pre-processing the interrupt before (and which is what appears to be more typical `JMP` style of `exec` in a chain). Sometimes the `JMP` style code is called a front end handler and the `CALL` style code is called a back end handler.

It is especially important that `INT21h` API 25h and 35h allow `exec` `INT21h` handlers to be hooked. This is a source of the trouble disassembling DOS, but it also makes it difficult for us to find the "Original `INT21h` Handler" calling `INT21h` API 35h (i.e. `exec` the `INT21h` handler linked list, that is, be a list of the most recently installed `INT21h` handlers). This might conceivably be the genuine `MSDOS\SYSTEM` handler but more likely it comes from `MSDOS\EXEC` or perhaps even some other valid DOS API `FUNCO032` or `IO$SERV` program from Chapter 2. `INT21h` API 35h simply returns the `next` of an interrupt chain. Finding the original `INT21h` or `INT21h` handler belonging to DOS usually requires finding the chain `next`. Usually, since there are a way because there might be back end handlers.

How can we find the actual `INT21h` and `INT21h` handlers provide the DOS itself when all we have is the address of the head of the `INT21h` or `INT21h` interrupt chain. There is a somewhat naive function that returns the `next` of an interrupt chain. And while there is a `next` which invokes DOS function `INT21h` API 1203h to return the DOS data segment, there is a `next` which invokes that return to the DOS code segment which remember may well be the HIMM.

One solution would of course be to boot on an absolutely bare bones system (I hope that `INT21h` and `INT21h` point to the original `MSDOS` handlers, thereby bypassing the whole process

can or how to follow interrupt chains. Or you could write a device driver to keep track of interrupts and send it very early in DOS initialization. But this is ridiculous! Clearly there must be some way to find out interrupt priority as the processor does this many times a second.

Unfortunately there is no standard mechanism for interrupt chaining. IBM and Microsoft at one point attempted to specify one for this purpose. David Thacker described in detail in *Microsoft System Journal*, July 1992, pp. 24-28, but unfortunately no one seems to use it. Rob Brown has provided a "INI 21h patch v. 4" described in the Interrupt List on disk to combat the extremely long interrupt chains that can result from plug a INI 21h, but again you can't rely on programs to do the right thing and use this protocol.

## Tracing a DOS INT 21h Call

It turns out that Microsoft provides with every copy of DOS an almost perfect solution to the problem of finding the actual DOS INT 21h and INT 21h handlers. This solution is none other than DBL G.

As you've noticed, DBL G has a command to assemble instructions in hex and a *t* key to step through instructions, as opposed to stepping over instructions. Even better, unlike some other software, step through debugger, the command in DBL G to run trace *through* an *IN*<sub>x</sub> instruction (or any other instruction) only works if DBL G does not treat *IN* as an atomic operation.

```
C:\MND0C2\CHAP6>intvec 21
INT 21h 0F95:32B6 MSCDEX
```

```
C:\MND0C2\CHAP6>debug
```

```

:
1985 0100 mov ah, 62
1985 0102 int 21
1985 0106 ret
985 0105
+
```

```

AX:6200 BX:0000 CX:0000 DX:0000 SP:FEEF BP:0000 SI:0000 DI:0000
DS:1985 ES:1985 SS:1985 CS:1985 IP:0102  NV JP EI PL NZ NA PD NC
1985 0102 CB21          INT 21
+
```

```

AX:6200 BX:0000 CX:0000 DX:0000 SP:FEEB BP:0000 SI:0000 DI:0000
DS:1985 ES:1985 SS:1985 CS:0195 IP:32B6  NV JP DI PL NZ NA PD NC
0195:32B6 80FC60          CMP AH,60
```

Notice that pressing *t* at the INT 21h instruction took us into the first line of the handler at 0195:32B6, not the next instruction at 0195:0104. This is exactly what one might expect from a processor that handles *push* as because of the way the single step interrupt works on the processor (see NICHAN's "Listing 1's" later in this chapter). Most debuggers don't allow hex and assembly, but every copy of DOS, since the beginning, does.

As you've noted, it's ironic to follow the INT 21h or INT 21h chain down into the bowels of DOS itself. Track Microsoft's step key tracing, either by continuously pressing *t* or by telling DBL G to trace out such as *trace* to trace a certain number of instructions, until the segment offset points to DBL G and a RET instruction, which in the example above is located at 985:0104. Now we're looking at the actual DOS INT 21h or INT 21h handler.

However, as a wise reader might have noticed right now, before we go any farther, that using DBL G to trace into INT 21h "won't work" because DBL G itself sees DOS and DOS, as we all know, is a recursive thing. It is absolute truth, a debugger that does not use DOS, such as Nu Mega's Soft ICE, is required since DBL G to tracing through DOS.

However, there are a handful of DOS functions that are reentrant, at least for the purposes of tracing with DEBUG. By examining the DOS code for INT 21h, we will soon see precisely what this reentrancy or lack thereof means. In the meantime, simply take it on faith that the DOS INT 21h functions shown in Table 6-1 are, with an important caveat that we'll get to in a moment, and thus traceable using DEBUG. SYMDBG is another debugger that uses DOS. With the exception of the undocumented INT 21h AH=64h, note that these are among the INT 21h functions that Microsoft's *MS-DOS 5.0 Internals Reference* Chapter 7 lists as callable from a critical error handler.

**Table 6-1: Reentrant MS-DOS Functions**

INT 21h AH=33h	Get/Set Ctrl-Break Get/Bios Drive Get DOS Version
INT 21h AH=51h and AH=62h	(Get PSP)
INT 21h AH=50h	(Set PSP)
INT 21h AH=64h	Set Drive Lookahead Flag

*Note: These functions are in DOS 3.00+ and the 4.20 line is disabled.*

It is desirable for MS-DOS to step out to Get/Set PSP functions for special treatment because this means that interrupt handlers can freely call these process manipulation functions (see Chapter 9 on ISRs—BIOS interrupt assembly language functions 53h and 64h merit this special attention). It would seem that the functions such as AH=25h and AH=39h to get and set interrupt vectors, might be more useful. On the other hand, depending on how 33h here means the interrupt handlers can't reset and set the DOS BREAK flag.

Let us now use DEBUG to examine and call a couple of these functions: INT 21h AH=62h (Get PSP) and 50h (Set PSP), which occurs when the function is called under DOS 6.0 in a configuration with a few standard DOS ISRs such as MMIOEN and DOSKEY. The documentation states that function 62h takes a parameter *dx* that the number 62h in AH and the *dx* function returns the current PSP value. You can probably guess that the DOS implementation of this function is rather simple, doing little more than copying 25h from the C:\BIOS\SPBIO.SYS DOS 6.0 segment. This location corresponds to the 0x00000000 of the Swapgate File Vec 500h, see FN\_21h AX=5100h in the appendix. However, you will see the processor execute most of code before DOS can do all gets and sets of entries in the other rows, so we skip the PSP operations.

As you can see, the key to the DEBUG provides here is that unlike SYMDBG, it can report traces into the *Next* instruction (as Figure 6-2 comments have been added to the listing). If DEBUG output, using *...* to make them stand out.

**Figure 6-2: Starting to Trace a Call to INT 21h AH=62h**

```
C:\MSDOS2\CHAP6>debug
>
1985:0100 mov ah, 62
1985:0102 int 21
1985:0104 ret
1985:0105
<
AX=6200 BX=0000 CX=0000 DX=0000 SP=FFEE BP=0000 SI=0000 DI=0000
DS=1985 ES=1985 SS=1985 CS=1985 IP=0102  NV UP EI PL NZ NA PO NC
1985:0102 C621          INT 21
<
; ; We have to keep tracing until the segment offset comes back to
; ; our own code, the RET instruction at 1985:0104
AX=6200 BX=0000 CX=0000 DX=0000 SP=FFEB BP=0000 SI=0000 DI=0000
DS=1985 ES=1985 SS=1985 CS=DF93 IP=32B6  NV UP DI PL NZ NA PO NC
DF93:32B6 BDFC60          JMP AH,60
```

```

-2

.. Running MEM /D showed that above is MSCDEX. This is consistent
.. with output from INTVECT program. Apparently MSCDEX is interested
.. in the undocumented DOS INT 21h AH=60h (Truename) function. Note that
.. we were running MSCDEX /S (for network sharing., usually MSCDEX doesn't
.. care about the INT 21h AH=60h call

AX=6200 BX=0000 CX=0000 DX=0000 SP=FFEB BP=0000 SI=0000 DI=0000
DS=1985 ES=1985 SS=1985 CS=0F93 IP=32B9  NV UP DI PL NZ NA PO NC
0F93.32B9 7405          JZ 32C0
-t

AX 6200 BX 0000 CX 0000 DX 0000 SP FFEB BP 0000 SI 0000 DI 0000
DS 1985 ES 1985 SS 1985 CS 0F93 IP 32B8  NV UP DI PL NZ NA PO NC
0F93 32B8 2E          CS:
0F93 32BC FF2+0232    JRP FAR [32B2]          CS:32B2+15FA
-t

.. MSCDEX decided it's not interested in our call to 21/62, so it chains
.. to the previous handler whose address it earlier retrieved (by
.. calling 21/35) and saved away (apparently in CS:32B2) before installing
.. (with 21/25) its own INT 21h handler.

AX 6200 BX 0000 CX 0000 DX 0000 SP FFEB BP 0000 SI 0000 DI 0000
DS 1985 ES 1985 SS 1985 CS 07F9 IP 15FA  NV UP DI PL NZ NA PO NC
07F9-15FA 80FC3F      CMP AH,3F

```

```

.. We're now in the previous INT 21h handler. MEM /D shows that
.. 07F9 15FA is SMARTDRV. Here it's (reasonably enough) interested in
.. whether we've called INT 21h AH=1Bh to read from a file (SMARTDRV
.. wants to see if the data we want from the file is actually already
.. in its cache.) But we called 21/62 not 21/3F so...

```

Well, you get the idea. Running DEBUG in this way is a bit tedious, and saving its output to a file is a better idea. As an experiment, I've saved debug output with an ipif script, such as 2162.SCR in Listing 6-3, and selected its output to follow. For a complete discussion of DEBUG scripts, see *PC Magazine* (APR 1987, *Free* 2nd edition, Chapter 9). Furthermore, rather than repeatedly hitting *t* to trace single steps through instructions, you can give the *trace* command a numeric parameter, for example *t 10* (or *t 32*) to trace a series of instructions.

### Listing 6-3. 2162.SCR

```

C:\UNDOC2\CHAP6>type 2162.scr
e
mov ah, 62
int 21
ret
, blank line below is crucial to leave assembly mode!

t 100
q

```

I've only paid attention to guessing how many instructions to trace. If you ask DEBUG to trace too far, it starts executing garbage. You only want to trace until you return to the RETI instruction you were called or at least not trace past it. The best bet is first try *t 10*, examine DEBUG's output to see if the traced instruction comes back, then try *t 57*, examine the output again, and so on. In any case, *t 100* happens not to work here; a larger number would be needed on machines with more ISRs that hook INT 21h installed.

Figure 6-3 shows a complete trace into an INT 21h AH=62h call from the time we issued the INT 21h until DOS returns to us with the current PSP in BX. Normally, all that you see (or want to

sect's) of an INT 21h call's your input and its output. But Figure 6-3 views the DOS call through the looking glass of assembly. Instead of looking down at DOS, you'll be inside DOS looking up at the INT 21h call. If you're slightly disoriented at first, but if you carefully study Figure 6-3 to the end, you'd have a much better understanding of what DOS is all about.

### Figure 6-3: Tracing a Call to INT 21h AH=62h

```
C:\WINDOWS\CHAP6>debug<2162.scr>2162.out
```

```
C:\WINDOWS\CHAP6>type 2162.out
```

```
-a
1985.0100 mov ah, 62
1985.0102 int 21
1985.0104 ret
1985.0105
-t 106

AX=6200 BX=0000 CX=0000 DX=0000 SP=FFEE BP=0000 SI=0000 DI=0000
DS=1985 ES=1985 SS=1985 CS=1985 IP=0102 NV UP EI PL NZ NA PO NC
1985:0102 C021 INT 21

AX=6200 BX=0000 CX=0000 DX=0000 SP=FFE8 BP=0000 SI=0000 DI=0000
DS=1985 ES=1985 SS=1985 CS=0193 IP=32B6 NV UP DI PL NZ NA PO NC
0193:32B6 80FC60 CMP AH,60

;;; As before (Figure 6-2), we're in MSCDEX /B now

AX=6200 BX=0000 CX=0000 DX=0000 SP=FFE8 BP=0000 SI=0000 DI=0000
DS=1985 ES=1985 SS=1985 CS=0193 IP=32B9 NV UP DI PL NZ NA PO NC
0193:32B9 7405 JZ 32C0

;;; The AX=XXXX BX=XXXX etc. dump that DEBUG shows each time usually
;;; isn't important here, so from now on we'll omit it and blank lines,
;;; except when the register dump is useful

0193:32BB 2E CS,
0193:32BC FF2E0232 JMP FAR [32B2] CS:32B2+15FA
0193:15FA 80FC3F CMP AH,3F

;;; As before, we're in SMARTDRV now

0193:15FD 7414 JZ 1613
0193:15FF 80FC00 CMP AH,00
0193:1602 7426 JZ 162A
0193:1604 301325 CMP AX,2515
0193:1607 7431 JZ 165A
0193:1609 80FC68 CMP AH,68
0193:160C 7442 JZ 1650

;;; Above provides a catalog of the DOS INT 21h function calls that
;;; SMARTDRV cares about: 3FB (read file), 6DB (disk reset), 2513h
;;; (set INT 13h vector), 68h (commit file). All this makes sense.
;;; For example, SMARTDRV uses 21/0D as a signal to flush the cache.
;;; For some calls, such as 21/0D, SMARTDRV doesn't JMP to the previous
;;; handler; instead, it does a far CALL and examines the 21/0D as
;;; the way back

0193:160E 2E CS,
0193:160F FF2E1423 JMP FAR [2314] CS:2314+0800

;;; We called 21/62; SMARTDRV doesn't care, so SMARTDRV chains to
;;; previous handler, CS:0800, which SMARTDRV earlier got from
;;; calling 21/35 before installing its own 21 handler with 21/25, and
;;; which is stored in CS:2314
```

```

C801 0800 9C          PUSHF
... Maa running with DOS-UMB, so some INT 21h handlers are running
... in upper Memory Don't know who the owner of this is!

C801 0801 FB          STI
C801 0802 300258      CMP AX,5802
C801 0805 7413       JZ 081a
C801 0807 300358      CMP AX,5803
C801 080A 7431       JZ 083D
C801 080C 80FC31      CMP AH,31
C801 080F 7503       JNZ 0814
C801 0814 9B         POPF

.. We can see that this handler cares about calls to INT 21h functions
;: 5802h =Get UMB Link, 5803h (Get UMB Link), 31h (F8) Wonder why
;: Anyway we called 21:42; the handler isn't interested in that, so it
;: chains to the previous handler

C801 0815 2E         CS:
C801 0816 FF2ECC01    JMP FAR [01CE]          CS:01CE+0023
0255-0023 F8E052ECC    JMP CC2E 058E

.. DRV shows that seg 0255h is a block mode device driver for
;: B: through I: it is a low-memory stub for DoubleSpace, located in
;: High memory Stacker uses the same area; both have signatures at
;: 0255 0000 DRV also shows that CC2E 058E is DBLSPACE$!DoubleSpace!

CC2E 058E 9C          PUSHF
CC2E 058F FB          STI
CC2E 0590 FC          CLO
CC2E 0591 1E          PUSH DS
CC2E 0592 DE          PUSH CS
CC2E 0593 1F          POP DS
CC2E 0594 C606C20700  MOV BYTE PTR [07C2],00    DS:07C2+00
CC2E 0599 53          PUSH BX
CC2E 059A 8ADC         MOV BL,AH
CC2E 059C 80FB6C      CMP BL,6C
CC2E 059F 7759       JA 05FA
CC2E 05A1 32FF       XOR BH,BH
CC2E 05A3 8A9F1305   MOV BL,[BX+0513]          DS:0575+00
CC2E 05A7 FFA76005   JMP [BX+05B0]            DS:05B0+05FA

DoubleSpace is sufficiently tied into DOS that it uses a jump table to
;: store a handler for every DOS function. The table at CC2E:0513 holds
;: byte offsets into code at CC2E 0580. Most DOS functions (including
;: our 21:42 call) are just passed on. Examining the table with the FTAB
;: program from later in this chapter shows that DoubleSpace cares
;: about the following INT 21h functions: 00 0A 0B 1D 1E 1F, 25 31,
;: 34 39 3A 3E 41 43 4B 4C 56 57 5D, 68. We know this from
;: running "ftab.exe 0513 6d 0813" | grep -v 00". For example, it books
;: 21:25 because (like SMARTDRV) it wants to know whenever someone sets the
;: INT 13h (BIOS Disk) vector

CC2E 05FA 5B         POP BI
CC2E 05FB 1F         POP DS
CC2E 05FC 9B         POPF
CC2E 05FD 2E         CS
CC2E 05FE FF2E0005   JMP FAR [05D0]          CS:05D0+109E

;: Trivial handling for our 21:62 call. Pass it to previous handler .

0116:109E 90          NOP

```



;; KERN /D shows that 0116h is MS-DOS. Finally!

```
0116:109F 9D      HOP
0116:10A0 EBCC00   CALL    116F
```

;; Hmm, DOS is calling some subroutine (which we've traced into)

```
0116:116F 9C      PUSHF
0116:1170 1E      PUSH   DS
0116:1171 06      PUSH   ES
0116:1172 51      PUSH   CX
0116:1173 56      PUSH   SI
0116:1174 57      PUSH   DI
```

;; We need to see the registers for the next few instructions

;; Note what happens to DS and ES

```
AX=6200 BX=0000 CX=0000 DX=0000 SP=FF0A BP=0000 SI=0000 DI=0000
DS=19B5 ES=19B5 SS=19B5 CS=0116 IP=1175 NV UP DI NG NZ AC PE CY
0116:1175 2E      CS
0116:1176 C5366711   LD     SI,[1167]          CS:1167=0080
```

```
AX=6200 BX=0000 CX=0000 DX=0000 SP=FF0A BP=0000 SI=0080 DI=0000
DS=0000 ES=19B5 SS=19B5 CS=0116 IP=117A NV UP DI NG NZ AC PE CY
0116:117A 2E      CS
0116:117B 43E68111   ES DI,[1168]          CS:1168=0090
```

```
AX=6200 BX=0000 CX=0000 DX=0000 SP=FF0A BP=0000 SI=0080 DI=0090
DS=0000 ES=FFFF SS=19B5 CS=0116 IP=117F NV UP DI NG NZ AC PE CY
0116:117F 890400   MOV   CX,0004
```

```
0116:1182 FC      CLD
0116:1183 F3      REPZ
0116:1184 A7      CMPSW
0116:1185 7407   JB   118E
```

```
;; DOS has just compared 8 bytes (4 words) at DS:SI 0000 0080, and
;; ES:DI FFFF 0090. If they are identical, DOS jumps somewhere
;; What is this? This particular run of DEBUG was conducted with
;; DOS-HIGH. DOS is in the HMA which is only reachable when the
;; machine's A20 address line is enabled. DOS is comparing 0000 0080
;; and FFFF 0090 because if the 8 bytes at these two addresses are
;; identical, it assumes that memory addresses are wrapping around, and
;; therefore that A20 is off. DOS can't call routines in the HMA if A20
;; is off. Thus, even when DOS-HIGH there must be a low memory stub the
;; code at 0116:109E is that stub which ensures that A20 is enabled before
;; calling DOS in the HMA. Here A20 was already on (0000 0080 and
;; FFFF 0090 were different), but A20 has been off; we would
;; have jumped to the subroutine at 0116:118E whose job
;; is to enable A20 by calling KMS function 5 (Local Enable A20).
;; If that function call succeeds, DOS will jump back here just as if
;; A20 had been enabled all along. If that function call fails, we're
;; in big trouble. DOS uses INT 10h AX=00h to display "A20 Hardware
;; Error" and goes into a dynamic halt. We'll come back to this
;; code later. Right now, A20 is enabled so...
```

```
0116:1187 5F      POP   DI
0116:1188 5E      POP   SI
0116:1189 59      POP   CX
0116:118A 07      POP   ES
0116:118B 1F      POP   DS
0116:118C 9D      POPF
0116:118D C3      RET
0116:10A3 2E      CS
```

0116.10A4 FF2E6A10 JMP FAR [106A] CS:106A=40F8

;; The low memory stub for DOS knows it can jump to DOS in the 106A, and  
;; here we go:

FDCB 40F8 FA CLI

;; We are now in The One True INT 21h Handler That this is at  
;; FDCB 40F8 in this particular configuration is the one piece of  
;; information we re after here, because now we can go and disassemble  
;; rather than trace at that address. Static unassembly is  
;; generally easier than dynamic tracing. But let a see the thing  
;; through to learn exactly how 21/62 is handled.

FDCB 40F9 80FC6C CMP AH,6C  
FDCB 40FC 77B2 JA 40B0

;; Any INT 21h function > 6Ch is an error ("In DOS 7.0,  
;; the upper limit is 72h," writes one tech reviewer.)

FDCB:40FE 80FC33 CMP AH,33  
FDCB:41D1 7218 JB 411B

;; Any INT 21h function < 33h will be handled at FDCB:411B

FDCB:41D5 74A2 JZ 40A7

;; 21 33 is special. It is handled at FDCB 40A7 (in this configuration)

FDCB:41D5 80FC64 CMP AH,64  
FDCB 41D8 7711 JA 411B

;; Any INT 21h function > 64h will also be handled at FDCB 411B;  
;; seems like 411B is the handler for "normal" DOS calls.

FDCB:41DA 74B5 JZ 40C1

;; 21 64 is another special function handled here at FDCB:40C1

FDCB:41DC 80FC51 CMP AH,51  
FDCB 41D1 74A4 JZ 40B5  
FDCB:4113 80FC62 CMP AH,62  
FDCB:4114 749F JZ 40B5

;; Finally DOS sees our 21/62 call and will handle it by jumping to  
;; FDCB 40B5. Notice that the same code also handles calls to 21/51, which  
;; makes sense since the two functions are documented as being identical.

FDCB 40B5 1E PUSH DS  
FDCB 40B6 2E CS  
FDCB 40B7 BE'EE7'D MOV DS,[30E7] CS 30E7=0116

;; DOS DS 0116h is stored in a variable kept at CS 30E7. This is  
;; the segment where things like SysVars and SDA live. This value is  
;; also returned from 2F/170 (see appendix).

FDCB:40B8 8B'E30D MOV BX,[0330] DS 0330=1408

;; Believe it or not, the previous line is actually the Get PSP function!  
;; We know that DOS keeps the current PSP at SDA+10h. In this  
;; configuration 21/5004 (Get SDA) returns 0116 0320. The Get PSP  
;; function just moves the WORD at 0116 0320 into BX. In other words,  
;; 21/62 (and 21/51) just return the WORD from SDA+10h. Dub.

FDCB:40BF 1F POP DS

```

FDCB:40C0 CF          IRET

;; DOS IRETs back to our code running in DEBUG
1985:0104 C5          RET

;; This is the RET statement in our DEBUG script
1985:0000 CD20        INT 20
0116:1094 90          NOP

;; Our script has already returned to DEBUG, which did an INT 20h return
;; to DOS. At this point, we start tracing all sorts of things we don't
;; care about. If we trace too far, we start to make DEBUG execute
;; garbage, which can hang the machine

```

-q

The most noticeable feature of the INT 21h trace in Figure 6-3 is the way that DOS extensions such as SMARTDRV and MSDIEN become indistinguishable from DOS itself. If any non-Microsoft DOS extensions such as Novell NetWare or Stacker had been running, they too would have appeared in the INT 21h chain, looking not a bit different from any of the Microsoft-provided software in the chain. The walk through the INT 21h chain in Figure 6-3 thus presents an excellent illustration of what DOS really is.

### Unassembling the Get/Set PSP Functions

As you can see, under normal circumstances with a few IIRs loaded, you have to wade through a lot just to get to the segment of code that actually performs the DOS Get PSP function. It should now be clear why INT 21h is called an "interrupt chain." As you'll see later, the INT 21h chain is typically much longer than the INT 21h chain, given the overhead of INT 21h on a typical machine. Programmers might even consider writing their own Get PSP calls to bypass this long interrupt chain. See how DOS implements Get PSP, when it eventually gets there — you can also see how to implement your own.

```

// Uses get_ada() from GETSDA.C (Listing 3-4a)
WORD my_get_psp(void)
{
    static WORD far *psp_ptr = (WORD far *) 0;
    if (! psp_ptr) // one-time init
        psp_ptr = (WORD far *) (get_ada() + 0x7D);
    return *psp_ptr;
}

```

Of course, this would eat out any IIRs or drivers that might actually need to see and respond to DOS Get PSP calls.

Having already seen the code that handles the Get PSP function, INT 21h AH=51h and 62h, we might as well also examine the code for Set PSP shown in Figure 6-4, though we can guess what it's going to look like — see later in Figure 6-4 where the 4039h address comes from.

**Figure 6-4: Implementation of INT 21h AH=50h (Set PSP) in MS-DOS 6.0**

```

-> 4 fdc8 40a9
FDCB:40A9 1E          PUSH    DS           ; save caller's DS
FDCB:40AA 2E          CS
FDCB:40AB 8E1EE73D     MOV    DS,[3DE7]     ; switch to DOS DS
FDCB:40AF 891E3003     MOV    [0330],BX     ; put caller's BX into CURR_PSP
FDCB:40B3 1F          POP    DS           ; restore caller's DS
FDCB:40B4 CF          IRET                ; done!

```

Other users, the Get and Set PSP functions just manipulate this word at offset 330h in the DOS vfat segment (offset 10h in the SFA). This provides a small taste of how DOS internally uses subfunctions as well as structures as SAs and the SFA. Thus:

```
void my_set_psp(WORD psp)
{
    static WORD far *psp_ptr = (WORD far *) 0;
    if (! psp_ptr) // one-time init
        psp_ptr = (WORD far *) (get_sda() + 0x10),
        *psp_ptr = psp;
}
```

This is provided merely as an example of what DOS itself does. Calling this function instead of INT 21h will not work because it bypasses SFA and other INTs that need to see Set PSP calls.

### Unassembling INT 21h AH=33h

A look at the assembly of the DEBUG output in Figure 6-3 shows that MS-DOS special-cases a label at offset 4052 (hex 1052) and just shows it in Figure 6-5. These functions are similar to the ones in DOS's function list in Table 6-1. While we will not quote a position to the assembly of the assembly function different from all other DOS functions, we do at any rate show the assembly address of the unassemble. Recall that this was our goal in tracing through DOS.

For example, INT 21h AH=33h is an on-disk function with a number of subfunctions relating to CD-ROMs. In the DOS vfat, the DOS version. For example, setting BREQ=ON ends up calling INT 21h AX=0000 with DI=1. In the accompanying code, at EBX=4052 handles this function:

```
FD0B:40FE 8DFE33      CMP AH,33
FD0B:4101 7218      JB 411B
FD0B:4105 74A2      JZ 40A7
```

We will not trace the entire trace of this function using the BREQ or any other DOS debug patch or the trace has been added to the output in Figure 6-5, which has also been scanned up slightly.

### Figure 6-5. Implementation of INT 21h AH=33h in MS-DOS 6.0

```
C:\ND002\CHAP6>debug
> r 3.0 a 47
FD0B:40A7 EB49      JRP 4052
> r 3.0 a 52
FD0B:4052 7436      CMP AL,06 ; functions 3300h through 3306h
FD0B:4054 7603      JBE 4059
FD0B:4056 8DFF      MOV AL,FF ; error: subfunction number too high
FD0B:4058 CF      IRET
FD0B:4059 5E      PUSH DS ; save caller's DS
FD0B:405A 2E      ESC
FD0B:405B 8A1E75D0  MOV DS,30E7 ; switch to DOS's DS, hrm, not truly
FD0B:405E 5D      PUSH AX ; reentrant after all!
FD0B:4060 56      PUSH SI
FD0B:4061 B13F05  MOV SI,0337 ; offset of break flag: SDA+17h
FD0B:4064 32E4      XOR AH,AH ; see if subfunc 0
FD0B:4066 0BC0      OR AX,AX
FD0B:4068 7504      JNZ 406E
FD0B:406A 8A16      MOV DL,[SI] ; 21/3300 -- get break flag
FD0B:406C EB35      JNP 40A3
FD0B:406E 48      DEC AX ; see if subfunc 1
FD0B:406F 7507      JNZ 4078
FD0B:4071 80E2D1  AND DL,01
FD0B:4074 B816      MOV [SI],DL ; 21/3301 set break flag
FD0B:4076 EB28      JMP 40A3
FD0B:4078 48      DEC AX ; see if subfunc 2
FD0B:4079 7507      JNZ 40B2
```

```

FDCB:4078 80E201      AND DL,01
FDCB:407E 8614      XCHG DL,[SI]      ; 21/3302 (IMBDC) get/set brk flag
FDCB:4080 E021      JMP 40A3          ; as single atomic operation XCHG
FDCB:4082 300300     CMP AX,0003      ; see if subfunc 3 (already subtracted 2)
FDCB:4085 7506      JNZ 408B
FDCB:4087 8A766900   MOV DL,[0069]    ; 21/3305 -- get startup drive
FDCB:408B EB16      JMP 40A3
FDCB:408D 300400     CMP AX,0004      ; see if subfunc 6 (already subtracted 2)
FDCB:4090 7511      JNZ 40A3
FDCB:4092 000600     MOV BX,0006      ; 21/3306 -- MS-DOS version 6.0
FDCB:4095 B200      MOV DL,00
FDCB:4097 32F6      XOR DH,DH
FDCB:4099 803E11200  CMP BYTE PTR [1211],00 ; is DOS=HIGH?
FDCB:409E 7403      JZ 40A3
FDCB:40A0 80CE10     DEI DH,10        ; DOSIMHRA flag
FDCB:40A3 5E        POP SI           ; done restore caller's regs
FDCB:40A4 58        POP AX
FDCB:40A5 1F        POP DS
FDCB:40A6 CF        IRET           ; return to caller

```

In addition to showing how DOS happens to handle function 33h, the code in Figure 6-5 also provides much support information that can be used to understand the disasm by using other parts of MS-DOS. For example, Microsoft documents INT 21h AX=3300h as returning the DOSIMHRA flag (DI). This code in Figure 6-5 shows DOS using the byte at DOS[DS], 211h to set DI. The code at DOS[DS], 211h must be the DOS\_HRHI indicator. It's not important by itself, but it's important to help you understand other parts of the code. Another example, DOS[DS], 121h, you may know, but this is the DOSIMHRA flag.

Similarly, functions 3300h and 330100h know how to get and set the C-0-C flag. Figure 6-5 shows these functions manipulating the byte at offset 0537h in the DOS\_05 segment. This byte must then be the C-0-C flag. Our final example in Figure 6-5 is how DOS sets the flag. Finally, Microsoft documents INT 21h AX=3400h as returning the startup device ID, and the code in Figure 6-5 also shows DOS setting DI from DOS[DS], 0069h (therefore, somewhere else in the code, where you see DOS[DS], 0069h, you can translate this to STARTUP\_DRIVE).

### Examining the Low Memory Stub for DOS=HIGH

Another interesting location to examine is the function that DOS calls and runs stackably when DOS\_HRHI (0010h) and the A20 line is disabled. The processor's A20 address line accesses memory only once, negatively. It's used on 386 and later processors by the A20 controller to cause the address wrap-around on address 0. In DOS\_HRHI, the A20 is off, DOS must first cause A20 to be on, can reach its own HIMEM.sys memory capability, then DOS's code is sophisticated enough to figure out how to check A20, the first place where a reference to it appears in any MS-DOS\_HRHI. Earlier, Figure 6-3, you saw that was created at 0116:1181. Figure 6-6 shows what this function actually does.

**Figure 6-6: DOS Function Called When DOS=HIGH But A20 Is Off**

```

0116:118e
0116:118E 53          PUSH  BX
0116:118F 50          PUSH  AX
0116:1190 8C8D      MOV   AX,55
0116:1192 2E        CS
0116:1193 A38610   MOV   [1086],AX
0116:1196 2E        CS-
0116:1197 89268B10  MOV   [108B],SP ; save caller's stack
0116:1198 8CCB      MOV   AX,CS      ; switch to a DOS stack, hmm, not
0116:1199 8EDC      MOV   SS,AX      ; reentrant at all if A20 off
0116:119F 0C4007   MOV   SP,0740    ; SD#480#end of Crit Err Stack
0116:11A2 8405      MOV   AH,05     ; MS func 5 + local Enable A20

```

```

0116 1144 2E          ES
0116 1145 FF1E6311   CALL   FAR [1163] ; SMS address from 2F/4310
0116 1149 0BC0       OR     AX,AX
0116 114B 740F       JZ     11BC ; failed; can't turn A20 on!

... jkey
0116 114D 2E          CS:
0116 114E A1261D     MOV   AX,[11086]
0116 1151 8E00       MOV   SS,AX
0116 1153 2E          CS:
0116 1154 8B268B10   MOV   SP,[11088] ; switch back to caller's stack
0116 1158 5B        POP   AX
0116 1159 5B        POP   BX
0116 115A EBC8       JMP   11B7 ; jump back into normal code (fig 6-3)
                                ; as if A20 had been enabled all along

... f41
0116 115C 840F       MOV   AH,0F ; come here if couldn't enable A20
0116 115E CD10       INT   10 ; get video mode
0116 1160 3C07       CMP   AL,07
0116 1162 7406       JZ     116A
0116 1164 32E4       XOR   AH,AH
0116 1166 8002       MOV   AL,02 ; set normal text mode
0116 1168 CD10       INT   10
0116 116A 8405       MOV   AH,05
0116 116C 32C0       XOR   AL,AL ; set display page 0
0116 116E CD10       INT   10
0116 1170 BEB812     MOV   SI,12B8 ; 12B8 -> "\nA20 Hardware Error\n$"
0116 1173 0E        PUSH  CS
0116 1174 1F        POP   DS
0116 1175 FC        CLD
0116 1176 AC        LODSB
0116 1177 3C24       CMP   AL,24 ; look for '$'
0116 1179 7409       JZ     11E4
0116 117B 840E       MOV   AH,0E ; write in TTY mode (use BIOS
0116 117D 8B0700     MOV   BX,0007 ; since can't make DOS calls
0116 117F CD10       INT   10 ; here!)
0116 1181 7407       JMP   11D6
0116 1184 1B        STI
0116 1185 EB1D       JMP   11E4 ; tight little loop (INTs on)

d 116 12b8
0116 12B0          06 0A 41 32 3D 2D 48 61          A20 Ha
0116 12C0 72 64 77 61 72 65 2D 45-72 72 6F 72 0B 0A 24 36 rdware Error. $6

```

Notice we've seen that DOS leaves the A20 bit on. This reduces the overhead of keeping the BIOS code in the HMA. DOS probably doesn't have to do the low-memory stub in Figure 6-6 very often.

For a long time DOS in the HMA is guarded with this low-memory stub. It brings up an interesting question: What about *data* in the HMA? MS-DOS doesn't put internal data structures such as the Current Directory Structure (CDS) and Session File Tables (SFT) up in the HMA because this would make some more third-party applications that peek and poke those (especially internal structures and the session file tables) have to ensure that A20 is enabled. However, DOS does keep its BUFFERS in the HMA. If a program can enable BUFFERS (and Chapter 8, see Listing 8-8, accesses the DOS sector buffer) or if some other session of DOS has FILESHIGH or LAN\$DRIVEHIGH statements that use HMA, and one tech reviewer, the program would need to check and possibly reenable A20 just like DOS does and game 6-6. But since, immediately, you just seen an trivial DOS call will ensure that A20 is turned on, perhaps a program that accesses data in the HMA merely needs to preface that access with a call to DOS call DOS will take care of checking the A20 state and, if necessary, calling VM\$turn on 5 to enable A20. But any TSR could turn it off. How frequently should programs that access the HMA check the A20 state? How much of a problem is this? Is the extra few bytes gained

by putting data in the HMA with this kind of uncertainty. "Ouch! This makes my head hurt," says one of the tech reviewers.

### Examining the INT 21h Dispatch Function

Of all the addresses we found through tracing, the INT 21h call the most important set of DOS's INT 21h handler seen in Figure 6-3 at 11C8:4018. This is really the piece of information we wanted at long last to see exactly what happens during an INT 21h call: we can now disassemble at this address. By tracing in INT 21h:402a, we only saw those snippets that happen to get executed when calling the Get PSP instruction; we can't look at the entire function. However, page 67, the DOS INT 21h handler, this time we've used SYMDIFF and added some labels as well as comments. In Microsoft's source code, this all-important function located at MSDOSPASM is called COMMAND.

Figure 6-7: MS DOS 6.0 INT 21h Dispatch Function

```

-u fdc8:40f8
FDC8:40f8 fa          CLI                    ; disable interrupts
FDC8:40f9 80fc6c     CMP                    ;
FDC8:40fc 7702     JA 4080                ; invalid function number

; step 1
FDC8:40fe 80fc53     CMP AH,33
FDC8:4101 7218     JB 4110                ; normal DOS function
FDC8:4103 76a2     JZ 40a7                ; do 21/33 (fig. 6-5)
FDC8:4105 80fc64     CMP AH,64
FDC8:4108 7711     JA 4110                ; normal DOS function
FDC8:410a 7685     JZ 40c1                ; do 21/64
FDC8:410c 80fc51     CMP AH,51
FDC8:410f 76a4     JZ 4085                ; do Get PSP
FDC8:4111 80fc62     CMP AH,62
FDC8:4114 769f     JZ 4085                ; do Get PSP (51==62)
FDC8:4116 80fc50     CMP AH,50
FDC8:4119 768e     JZ 40a9                ; do Set PSP (fig. 6-6)

normal_DOS
; step 2
; caller's f.eqs, CS, and IP of course already pushed on the stack by INT
FDC8:411b 06          PUSH ES                ; Save regs on caller's stack
FDC8:411c 1e          PUSH DS                ; The order is important, as
FDC8:411d 55          PUSH BP                ; later on different INT 21h
FDC8:411e 57          PUSH DI                ; functions will access the
FDC8:411f 56          PUSH SI                ; caller's original registers
FDC8:4120 52          PUSH DX                ; by treating this stack frame
FDC8:4121 51          PUSH CX                ; as a structure. See 2f/3218
FDC8:4122 53          PUSH BX                ; for example, caller's BX
FDC8:4123 50          PUSH AX                ; is at offset 2, ES at 10h

; step 3
FDC8:4124 8c0b     MOV AX,DS
FDC8:4126 2e8e1ee730  MOV DS,CS [30e7]      ; get DOS DS
FDC8:4128 a3e005     MOV [D5EC],AX         ; save caller's DS
FDC8:412e 891eeaf05     MOV [D5EA],BX         ; save caller's BX
FDC8:4132 a18405     MOV AX,[0584]         ; SDA+264h = ptr to stack frame
FDC8:4135 a3f205     MOV [D5F2],AX         ; containing user registers
FDC8:4138 a18605     MOV AX,[0586]         ;
FDC8:413b a3f005     MOV [D5FD],AX

; step 4
FDC8:413e 33c0     XOR AX,AX              ; set AX=0
FDC8:4140 a27205     MOV [D572],AL
FDC8:4143 f606301c01  TEST Byte Ptr [1030],01 ; Is Min3 Enh running?
FDC8:4148 7503     JNZ 414b

```

UNDOCUMENTED DOS, Second Edition

```

; For using line only if Windows 3 Enhanced mode not running!
FDCB:414A A33E03      MOV     [033E],AX      ; set machine ID to zero

; step 5
FDCB:414D FE02103     INC     Byte Ptr [0327] ; increment INDOS flag

; step 6
FDCB:4151 89268405     MOV     [0584],SP      ; SDA+264h
FDCB:4155 8C168605     MOV     [0586],SS      ; save current stack ptr
FDCB:4159 A13003      MOV     AX,[0330]      ; get current PSP
FDCB:415C A13C03      MOV     [033C],AX     ; SDA+1Ch = SHARE, NET PSP
FDCB:415F 8E08      MOV     DS,AX          ; point DS at caller's PSP
FDCB:4161 5B        POP     AX
FDCB:4162 5B        PUSH   AX             ; get back caller's AX
FDCB:4165 89262E00     MOV     [002E],SP     ; save current stack ptr
FDCB:4167 8C163000     MOV     [0030],SS     ; in caller's PSP
FDCB:416B 2E8E16E73B   MOV     SS,CS:[3DE7]
      INT 21h AX 5000h (Server Function Call) jumps to here

FDCB:4170 DCAD07      MOV     SP,07AD       ; switch stack to 07ADh-SDA =
                        ; SDA+480h=end of Crit Err Stk

; step 7
FDCB:4175 1B        BTI                     ; reenable interrupts
FDCB:417A 8CB5      MOV     BX,SS
FDCB:417E 8E08      MOV     DS,BX         ; point DS at DOS_DS
FDCB:417F 91        XCHG   AX,BX         ; caller's AX into BX
FDCB:4179 13C6      XOR    AX,AX
FDCB:417B 36A2F605     MOV     SS [05F6],AL   ; extended open off?
FDCB:417F 36B12611060008 AND    word Ptr SS [0611],0800
FDCB:4186 36A25703     MOV     SS [0557],AL   ; set different vers to 0
FDCB:418A 36A24C03     MOV     SS [034C],AL
FDCB:418E 36A24AD3     MOV     SS [034A],AL
FDCB:4192 40        INC    AX
FDCB:4194 36A27803     MOV     SS [0358],AL   ; okay to do INT 28h

; step 8
FDCB:4197 93        XCHG   AX,BX         ; get back caller's AX
FDCB:4198 8AD6      MOV     BL,AX         ; DOS func num into BL
FDCB:419A D1E3      SHL    BX,1          ; make DOS func number into word ofs

; step 9
FDCB:419C FC        CLD
FDCB:419D 0AE4      OR     AH,AH
FDCB:419F 7417      JZ     41B8           ; AH=0 (terminate process)
FDCB:41A1 8D1C59     CMP    AH,59
; if 21 59 (get critical error), bypass code that turns off critical error!
FDCB:41A4 7444      JZ     41EA         ; AH=59h (get extended error)
FDCB:41A6 8D1E0C     CMP    AH,0C
FDCB:41A9 7700      JA     41B8         ; AH > Dch

INT21_D1_THRU_0C
; step 10
FDCB:41AB 36805F700300 CMP    Byte Ptr SS [0320],00 ; critical error set?
FDCB:41B1 7537      JNZ   41EA         ; if so, stay with crit error stack
FDCB:41B3 BCAD0A     MOV     SP,0AAD       ; SDA+780h=end of Char I/O Stack
FDCB:41B6 EB32      JMP    41EA

INT21_0D
INT21_ABOVE_0C ;, except (normally) 33h, 50h, 51h, 59h, 62h, 64h
; step 11
FDCB:41B8 36A33A03     MOV     SS [033A],AX
FDCB:41BC 36C606270301 MOV    Byte Ptr SS [0323],01 ; crit err locus
FDCB:41C2 36C606200300 MOV    Byte Ptr SS [0320],00 ; turn off crit error
FDCB:41C8 36C6062203FF MOV    Byte Ptr SS [0322],FF ; crit err drive#

```



```

; Windows Enhanced mode patches next four lines into a far call.
FDCB:41CF 50          PUSH  AX
FDCB:41CF 0482       MOV   AH,82
FDCB:41D1 0D2A       INT   2A          ; End crit section
FDCB:41D3 58          POP   AX

FDCB:41D4 36C06580300  MOV   Byte Ptr SS [0358],00 ; no INT 2Ah
FDCB:41D8 BC2D09       MOV   SP,D920      ; 50A+600h = end of Disk Stack
FDCB:41D0 36F6063703FF  TEST  Byte Ptr SS [0337],FF ; 50A+17h break flag
FDCB:41E3 7405       JZ    41EA
FDCB:41E5 5D          PUSH  AX          ; BREAK-ON, so
FDCB:41E6 EB964E       CALL  907F        ; check ctrl-break
FDCB:41E9 58          POP   AX

; step 12
;; Next four lines are the key call through dispatch table
;; BX holds caller's INT 21h function number SHL 1 word offset
FDCB:41EA 2E809F9E3E     MOV   BX,C5 [BX+3E9E] ; get func handler addr
FDCB:41EF 3687FEEA05     XCHG  BX,SS [05EA]   ; move func ptr into var
FDCB:41F4 368E1EEC05     MOV   DS,SS [05EC]  ; switch to caller's saved DS
FDCB:41F9 36FF16EA05     CALL  SS-[05EA]     ; call func handler addr
;; We've just called the DOS function for the specific DOS function in AH

; step 13
;; Now into cleanup preparatory to returning to caller
FDCB:41FE 3680768A00FB  AND   Byte Ptr SS [0086],FB
FDCB:4204 FA          CLI
FDCB:4205 2E8E1EE73D     MOV   DS,C5 [3DE7]  ; switch back to DOS DS
FDCB:420A 8D5E85D000     CMP   Byte Ptr [0085],00
FDCB:420F 7527       JNZ   4218
FDCB:4213 FE0E2105     DEC   Byte Ptr [0321] ; decrement InDOS
FDCB:4215 8E168605     MOV   SS,[0586]    ; switch back to caller's
FDCB:4219 8B2B8405     MOV   SP,[0584]    ; stack
FDCB:421D 80EC       MOV   BP,SP
FDCB:421F 8B4600       MOV   [BP+00],AL
FDCB:4222 A1F205       MOV   AX,[05F2]
FDCB:4225 A38405     MOV   [0584],AX    ; caller's SP
FDCB:4228 A1F005     MOV   AX,[05F0]
FDCB:422B A38605     MOV   [0586],AX    ; caller's SS
FDCB:422E 58          POP   AX          ; put back caller's
FDCB:422F 5B          POP   BX          ; registers, including
FDCB:4230 59          POP   CX          ; any changes the DOS
FDCB:4231 5A          POP   DX          ; function made to them
FDCB:4232 5E          POP   SI
FDCB:4233 5F          POP   DI
FDCB:4234 5D          POP   BP
FDCB:4235 1F          POP   DS
FDCB:4236 07          POP   ES
FDCB:4237 CF          RET

```

The dispatch table in Figure 6-7 is the heart of DOS. It is executed every time a program issues an INT 21h DOS dispatch function or the DOS equivalent of the function issued on UNIX, which has no equivalent in systems such as Rich's *Plan of the UNIX Operating System* [4-165, 168] and Antognini's *UNIX System Environment* [pp. 21-23]. The discussions of several other UNIX books provide a useful background for understanding the INT 21h dispatch code. However, UNIX is discussed separately between applications and the operating system. The discussion of system calls and the transition from user mode to kernel mode. Again, since there is something like a transition from DOS through DOS extensions such as Windows to a separation between the application and the protected mode and DOS running in real mode, we call this an important separation. DOS usually switches from the application's stack to the DOS own. This important aspect of DOS will be discussed in detail below.

Near the top of the function, commented "step 1", you see how DOS picks off a handful of special error codes (3h, 4h, 51, 62h and 50h). These of course are none other than what we've been calling the recurrent DOS functions. Here, "recurrent" simply means that while the above code is executing, after it has passed one unit of CUI, and before it has executed the closing INT, it could be interrupted by an interrupt handler, and the interrupt handler could call one of these five functions. These five functions are recurrent simply in the sense that DOS handles them without switching stacks and incrementing the InDOS flag. Thus an interrupt handler can call these functions, even if the InDOS or critical error flag is set.

As a large section of course, these functions aren't really recurrent given the way that, for example, the Set PSP function writes to a global variable (see Figure 6-4). MS-DOS's extensive reliance on global variables makes a complete, non-recurrent, Furthermore, InDOS=HIGH and the A20 gate is off. DOS registers to be shown has to switch stacks. But in any case, it should not be clear why we make INT 21: AH=02 contrast with DOS 4s and not say INT 21h AH=52h, DOS handles the after function only after switching stacks and incrementing InDOS.

Next, step 2, the INT 21 dispatch code pushes the caller's registers on to the caller's stack. The caller's of course, simply in whatever program issued the INT 21h call. This can be slightly disorienting because of course we're used to thinking about INT 21h from the caller's perspective and now we're looking at it from DOS's point of view. These pushed registers form a structure that many DOS functions exist on. Under documented INT 21: AX=4218, Get Caller's Registers, see appendix, returns a pointer to this structure.

At step 3 (Figure 6-7), DOS saves away the caller's DS and BX again and switches from the caller's DS to its own DS. DOS keeps DS in a variable, accessible through DOS's CS. It is also available on calling INT 21: AX=1205h, see get dos ds. Listing 6-2. Note that even though DOS's INT 21: the DOS code is in the HMA, the data segment is still in low memory. This is necessary because many existing DOS programs rely on the ability to reach DOS internal data structures and write to them. They check the status of the A20 gate. Microsoft has to introduce improvements such as DOS 5.00 HLL or other like existing applications.

The next step in the code does step 4, check a variable at 1030h to see whether Windows 3.x is in enhanced mode or Windows 386 2.x is running. Since most of us think of Windows as something that is "on top of" DOS, it is a bit disconcerting at first to learn that DOS 5.0 and higher can run under Windows. As discussed in Chapter 1, however, this part of the intricate DOS Windows interface is implemented using documented functionality. In its INT 21h handler, MSDOS5\SYSTEM\EXPL\AX=1605h Windows in Enhanced and AX=1606h exit broadcasts, the code for AX=1605h will set a variable at 0300h, set also set the byte at 1031h, and the code for AX=1606h will set this variable to zero as a kind of Low Windows flag. It's important to underline that this variable named mode only. DOS doesn't care one way or the other about Standard mode.

If Windows Enhanced mode is not running, then DOS zeroes out a variable at 0331h. VM ID is used by DOS is the machine ID. If Windows Enhanced mode is running, the DOS\MSGR\VMID is available in Chapter 1, has smacked a virtual machine ID in here. DOS uses this VM ID to manage SEEs.

Next, step 5, the code increments the InDOS flag, which is simply a variable at 0321h (SDA+1) in the DOS Ldt segment. The undocumented function INT 21h AH=34h (Get InDOS Flag Address) returns a pointer to this variable.

If the InDOS flag has been set, so we're now "in DOS". Of course, we were in DOS before, but the significant fact of this spot is that DOS is about to switch stacks. Switching stacks requires a guard or semaphore, and that's the InDOS flag. Notice, however, that while DOS increments the InDOS flag, it does not check it before proceeding. Thus InDOS is not a true semaphore. If the processor is interrupted in the middle of this code, or rather a little further on when DOS reenables interrupts with an SUI instruction, the code can be reentered.

In other words, DOS does nothing to enforce its requirement that only one caller at a time execute inside the INT 21h code. Obeying the InDOS flag is merely a courtesy, not that it is vital that programs do, because this convention, by not so making an INT 21h call a non-InDOS, is not almost always cause problems — or one thing DOS relies on many global variables. If, for example, DOS were working with a particular hard disk cluster to serve an INT 21h file I/O call, and an interrupt handler that ignored the InDOS flag made a file I/O call to DOS before DOS had finished with the first call, DOS would probably use the second caller's cluster to satisfy it, not the first caller's request. Good variables do not work any a last-in, first-out stack. It is vital that interrupt handlers check the DOS sector, using INT 21h calls. No one did, I take Microsoft's word to document InDOS and the INT 21h AH=34h function that returns a pointer to it.

Ignoring InDOS calls cause another problem. Because the code at step 5 in Figure 6-7 increments InDOS, executing DOS means that InDOS will take on a value of two or greater. This is bad because the internal DOS function that checks for Ctrl-C only does so when CMI, Bty, Ptn, IS, DOS, OI, I, or InDOS is 2 or greater. DOS won't check for Ctrl-C except IRR, AK, OS.

There is a catch to why InDOS can be safely ignored. If the state DOS state, including all three DOS stacks, is saved and restored by each caller, and if each such caller observes the DOS critical sections by looking INT 21h. The SDV ISR technique put forward in Chapter 9 is an approximation of this method, though it is an approximation because the SDV does not include the entire DOS state.

Returning to step 6 in Figure 6-7, you can see the beginnings of the stack switching code. How does DOS switch away from the user's stack to one of its own? First, it saves away the caller's current SS:SP. Next, DOS gets the current PSP (at 01330h) or SDV-40h, and uses it to save the caller's SS:SP at offset 21h in the caller's PSE. Finally, it sets SS:SP to a DOS stack. Depending on the DOS function number, it may switch again to differ on DOS stack.

What is the purpose of stack switching? Why not just use the caller's stack? Wouldn't that make DOS much more efficient? Yes, it would. It would make an INT 21h call actually only 18h bytes on the caller's stack (see Table 6-2). If the user could be forced on to provide a large enough stack, DOS could even save the stack. Unfortunately, DOS has to accommodate programs with unknown stack sizes. This complicates DOS's operations, and helps make DOS inefficient.

At the very end of step 6, where DOS points SP at the Critical Error stack is a location called Redip in the source code, a vector table, named INT 21h AX=51000h. Service Function Ctrl-C jumps this vector on as a back door into the INT 21h dispatcher. It is a new stack-aware program books they call it can be used by one machine to distribute INT 21h calls on another machine, or perhaps to another Windows virtual machine.

Skipping over a burst of the code at step 7, which zeroes out several variables in the DOS data segment, we come to step 8, where the code takes the caller's AH, with the critical DOS function number, and inserts it into word offset 1A. This will be important later on.

Next, step 9, DOS examines the function number in AH. If AH=59h, Get Extended Error is being called. DOS proceeds directly to step 12, where the user's function 59h will be called. It stays on the Ctrl-C Error stack, bypassing more stack switching code on steps 10 and 11, and bypassing code that will later generate an error on pertaining to any pending Critical Error.

If one of the Ctrl-C Misused Call after I/O functions, INT 21h AH=1 through AH=0Ch, is being called, DOS, step 10, points SP at 04A0h, which is the top of the character I/O stack, located in the Swap area of a Area, see Appendix. However, if there is a pending critical error, DOS stays with the Critical Error stack and was successful. This is not surprising, since Microsoft documents these functions. *MS-DOS 5.0 Reference* is calculable from a critical error handler. Notice that DOS does not carry off critical error information for functions 1 through 0Ch. As you can see, much of the core DOS code accommodates critical errors.

Finally, after DOS function number 0 (Terminate Programs) or anything greater than 0Ch but not 59h, and for one of the special functions that were picked off earlier in step 1 and which DOS did not process on the caller's stack, DOS step 11 switches to the disk stack. Thus, there are three DOS stacks:

- **Critical Error** (an array), used for function 59h and for functions 1 through 0Ch when a critical error is pending, and used temporarily for any DOS function call (DOS-ERR) but A20 was.
- **Character I/O Stack**, used for functions 2 through 0Ch in the absence of a critical error.
- **Disk Stack**, used for everything else. Calling any of the special functions with the IN1 21h AX=5000h causes a function call also ends up using this stack, though in practice indirectly, by the special instruction 21-5100 crashes the machine.

In the absence of functions remaining on the disk stack, the code (step 11) carries out a number of tasks, including critical error handling (undocumented IN1 2Ah AH=82h) to end any critical sections (see below), and a kernel-mode Ctrl-C check (flag at SDA 47h). In Figure 6-5, you saw the code for IN1 2Ah AX=800h that sets a busy flag when a user types BREAK. DOS now wants to see where DOS was interrupted (by DOS-MAC) that is, to see if the flag at SDA 47h is now zero. DOS calls a subfunction (undocumented 901h) to check Ctrl-C for the functions that come through here. Otherwise, DOS does not check Ctrl-C for functions 1 through 0Ch. As noted earlier, the DOS internal function to check Ctrl-C will only do so if IN1 0Ch=1.

We saw earlier that IN1 2Ah AH=82h normally is the IN1 2Ah handler in DOS, does an immediate IO operation, no operation. However, other programs can take over IN1 2Ah and/or patch DOS. We saw earlier that some particular uses IN1 2Ah critical sections because it runs free memory management. DOS has a concept of a simple copy of MS-DOS. Because the INDOS flag is just a flag, the VMM can't see it, DOS has its own copy, it cannot be used by other access to DOS, so the VMM. Now you know where the INDOS flag is located, as different VMMs, and be in different parts of DOS it is not same. What different parts. Different critical sections can be set and stored in IN1 2Ah AH=80h and 81h, see appendix. DOS call to IN1 2Ah AH=82h is a signal to start a critical section. DOS, such as Windows or networking software, can restart any task VMM that was suspended because it was waiting for a critical section. For additional information on critical sections, see Chapter 1 and Chapter 9, see REFERENCE in Listing 9-19. Also see Microsoft's *MS-DOS Technical Reference*, p. 11, which briefly discusses critical sections in the context of the MRC specification.

As discussed in earlier chapters, the DOS(MR-VAD) Windows Enhanced mode patches the IN1 2Ah AH=80h and IN1 2Ah AH=81h to register into a table into Windows. When Windows exits, it will restore registers back to the original code.

When a program calls a program, Ctrl-Break, and/or a critical section, which do dominate the DOS system, the program must have a real-time goal, which is that a program wants to call a DOS function. A program of software, DOS accomplishes this main goal in only a few lines, while other situations will use instructions and some assembly code.

Having selected an appropriate stack, saved caller's registers, and so on, step 12 in Figure 6-5 is the simplest and the most important. Recall that step 8 moved the function number in AH into AX and then pushed two DOS words, so this value is an offset into an array of function pointers, the function DOS array. Hence, the offset is CS:SI\*91, so that for example the address for DOS function 0 is CS:SI\*91, function 1's address is CS:SI\*A0, function 2's address is CS:SI\*A2, and so on. Since this array holds two-byte words, not a pointer, you can't use it to hook individual DOS functions. All handlers must be located in a single segment (here, F008h). We come back to this area of function pointers momentarily; it is very important to us. In any case, having retrieved the

address of the handler for the DOS function being called. DOS calls the handler. In a bit more from the user wanted has now been called.

In step 13 you can see the appropriate handler for the DOS function. DOS decrements the INT05 flag, switches back to the caller's stack and pops back the caller's registers from the register image created in the stack back in step 2. As you'll see in a moment, the handler for the specific DOS function probably switches a register image. Finally, DOS returns to the caller via an INT3. Since INT3 pops the flags off the stack, the specific DOS functions have to set or clear the carry flag by modifying its image, but the processor pushed on the stack as part of the initial INT3 (see the comment to step 2).

Seeing the DOS dispatch code in Figure 6-7 it should now be clear why a debugger traces through an INT3's INT05h callback works but tracing, for example, through a call to INT3's AH=52h wouldn't. A call to AH=52h would involve switching stacks, messing with global variables in the DOS data segment, and so on. THE BUGSIN uses DOS, so you wouldn't set up instead of acting through one of the DOS calls that THE BUGSIN would be making to display our information. A simple mess! One alternative of course would be a debugger that bypasses DOS, such as SoftICE or Tim Patterson's *STRAPON* from the first edition of *Undocumented DOS*, which, however, did not support tracing through INT3 instructions.

### Examining the INT 21h Dispatch Table

We really do need to trace through INT 21h a bit more. We now have the address of COMMAND.COM's One True INT 21h Handler and the address of the instruction pointer array called Dispatch in the DOS source code (in EXE21.DOC) that we assembled at binary, but set the trace under processor, so to speak.

To find the code that handles each specific DOS function, we would do something more than dump out the Dispatch table, which you can see from step 12 in Figure 6-7's source at FD083191. It's a bit of a pain, but we can dump out the SYMDOS.BIN data instead:

```
-du fdc8:3e9e
FDCB:3E9E A1F6 54E0 54E9 559F 550C 55C2 541C 544B
FDCB:3EAE 51BA 5214 5220 55D6 55E0 4DA1 4C7B 5C0C
FDCB:3EBE 5688 5D0F 5E73 5625 50CB 5D00 50B1 56F9
FDCB:3EE4 44D0 4C73 4C68 4B2B 4D2F 44D0 44D0 4D71
FDCB:3EE4 44D0 5D05 5D0A 5639 56D0 4C9A 4EB6 5D6C
FDCB:3EE4 5D01 4D22 4839 4856 4876 48B7 4AA6 4C54
FDCB:3EFE 4A1C 4194 4D73 4D52 4D59 4CBA 4C2B 4CC9
FDCB:3FDE 444D 4D01 4D29 4D65 4FE6 4FDF 472A 4B39
2, etc
```

You can double-check that a given code file is looking for a known function. Let's see what the table shows for `4405` (62h), although we know it usually gets pushed off in step 1 of Figure 6-7, only coming through the table in a special case, even if an INT 21h AX=51000h executes function call 4AH=62h):

```
-du fdc8:43e9e+62*2
FDCB:3F62 40B5

-u fdc8:40b5
FDCB:40B5 7E          PUSH    DS
FDCB:40B6 2E8E1EE73D      MOV     DS,CS:[30E7]
FDCB:40BB 8D7E4003      MOV     BX,[0330]
FDCB:40BF 7F          POP     DS
FDCB:40C0 CF          IRET
```

That's it! So you now have the DOS dispatch table and can examine at will the code for a DOS function you're interested in.

Examination of this table and others like it is made easier with a short C program, FTAB.C, shown in Listing 6-4. FTAB can display tables of bytes, 1, words, 2, or dwords, 4.

#### Listing 6-4: FTAB.C

```

/* FTAB.C */
#include <stdlib.h>
#include <stdio.h>
#include <dos.h>

typedef unsigned char BYTE,
typedef unsigned short WORD,
typedef unsigned long DWORD,
void fail(const char *) { puts(a); exit(1); }
main(int argc, char *argv[])
{
    char *prefix;
    void far *tab,
    BYTE far *btab,
    WORD far *wtab,
    DWORD far *dtab;
    WORD seg, ofs;
    int num_func, size, i;

    if (argc < 4)
        fail( "usage: ftab <seg ofs> <num_func | ?> [prefix] [size]");
    sscanf(argv[1], "%04X:%04X", &seg, &ofs);
    tab = (void far *) MK_FP(seg, ofs);

    if (argv[2][0] == '?')
    {
        num_func = ((BYTE far *) tab), /* first BYTE is #func */
        tab = ((BYTE far *) tab + 1), /* then array of WORDs */
    }
    else
        sscanf(argv[2], "%04X", &num_func);

    prefix = (argc > 3) ? argv[3] : "func";
    size = (argc > 4) ? atoi(argv[4]) : 2, /* default to WORD table */
    switch (size)
    {
        case 1:
            for (i=0, btab=(BYTE far *)tab, i<num_func; i++, btab++)
                printf( "%02X%13s%02X\n",
                    *btab, prefix, i);
            break;
        case 2:
            for (i=0, wtab=(WORD far *)tab, i<num_func; i++, wtab++)
                printf( "%04X:%04X%13s%02X\n",
                    seg, *wtab, prefix, i);
            break;
        case 4:
            for (i=0, dtab=(DWORD far *)tab; i<num_func; i++, dtab++)
                printf( "%08X%13s%02X\n",
                    *dtab, prefix, i);
            break;
        default:
            fail("size only 1 (byte), 2 (word), 4 (dword)");
    }
    return 0;
}

```

To generate a list of the 6Dh-0 through 6Ch-different DOS INT 21h functions, the Jcrs\_72h.net 6Ch is the lowest function number in the DOS 7.0 component of Chkdsk7's source code reviewer (see <http://www.fhlab.com/available-at-FDK-831-91/>). Figure 6-8 shows sample output from FHAB.

**Figure 6-8: The INT 21h Dispatch Table Displayed by FHAB**

```
C:\UMBD0C2\CHAP6>ftab fdc8 3e9e 6d int21 2
FCB:41F6 int21_00
FCB:54E0 int21_01
FCB:54E9 int21_02
FCB:559F int21_03
; ...
FCB:4C9A int21_25
; ...
FCB:4059 int21_34
FCB:4CBA int21_35
; ...
FCB:A10F INT21_38
FCB:A72A INT21_3E
FCB:A839 INT21_3F
FCB:AB9F INT21_40
FCB:BD38 INT21_41
FCB:ABA4 INT21_42
; ...
FCB:40A9 int21_50
FCB:40B5 int21_51
FCB:40B5 int21_52
FCB:4DD6 int21_53
FCB:4AA1 int21_54
FCB:4EA5 int21_55
FCB:BD5E int21_56
FCB:A90C int21_57
FCB:4448 int21_58
FCB:4CD0 int21_59
FCB:BD19 int21_5A
FCB:BD01 int21_5B
FCB:BD98 int21_5C
FCB:A531 int21_5D
FCB:AAA9 int21_5E
FCB:A9AA int21_5F
FCB:AEA8 int21_60
FCB:44D0 int21_61
FCB:40B5 int21_62
; ... etc
```

Confirming that the table is correct, you can see that int21\_51 and int21\_62 are located at the same address, FDK 8-40B5, as they should be.

### Get SysVars and the Caller's Registers

To check that the Jcrs\_72h.net 68h correctly executes another function that should be simple, INT 21h Ah (32h) — `get_sysvars` — far pointer to SysVars in FS:BX. According to the FHAB output, the `get_sysvars` function (52h) should be at FDK 8-4D65, so you can use SYMDIFF or DBI-BI-Cross to access that far address. Figure 6-9 shows the results.

**Figure 6-9: MS-DOS 6.0 Implementation of INT 21h AH=52h (Get SysVars)**

```
> fdc8 4d65
FCB:4d65 EB1AF5 CALL 42B2
FCB:4d68 C764022600 MOV Word Ptr [SI+021,0026
FCB:4d6b 8C5410 MOV [SI+103],SI
FCB:4d70 C3
```

In fact, calling Get SvVars in this particular configuration does return 0116-0026, so the hard-wired 0026h entry looks correct. But what is going on here? How come we don't see \$S0026 being protected? ENBX. What are \$I+021 and \$I+10h? To answer these questions, let's examine the subroutine being called at 4282h:

```

mov     edx, 4282
fcb     4282, 2E8E1EE730      mov     dx, cs:[30E7]
fcb     4282, C5368405        lod     si, C0584
fcb     4280, C3              ret

```

CS:4030 is just our old 11 and, see step 3 in Figure 6-7, the DOS data segment whose value DOS keeps in its code segment. DOS stores the value of DS in its code segment because, when an INT occurs, DS points out, but CS is. So this subroutine is just setting itself up to use DOS's DS, just as CS is, of 11, aka. Figures 6-3, 6-4, and 6-7 for Get PSP, Set PSP, and the INT 21h dispatch.

The point of the calls DS=SI with something at DOS:584h (i.e., step 6) of the INT 21h dispatch code in Figure 6-7 is to have DOS set the word at DOS:584h to the caller's SSIP. In other words, DOS:584h contains a pointer to the caller's stack, with all the registers that were pushed onto it during step 2 of Figure 6-7, as well as, as part of the actual INT instruction. Sure enough, the contents of CS:4282 (our call to DOS:584h) in this configuration happens to be \$I+260h, which is the physical address of a pointer to the stack frame containing the user registers on entry to the INT 21h call.

So why, if it were at 110:8498? Calls DS=SI with a pointer to the caller's pushed register structure. Given the offset in which steps push registers, it won't surprise you to learn that the caller registers are pushed in the format shown in Table 6-2.

**Table 6-2: MS-DOS Caller's Register Structure**

00h	AX
02h	BX
04h	CX
06h	DX
08h	SI
0Ah	DI
0Ch	BP
0Eh	ES
10h	EB
12h	EP
14h	CS
16h	Flags

In Figure 6-9, the code for instruction 82h at 110:84D06h moves 26h into \$I+21 and DOS's SS (DS) into \$I+1E. DS=SI points to the caller's register structure, where offset 2 is BX and offset 10 is ES. So we're looking exactly what you'd expect of the caller's ENBX to DOS: DS:0026, the register contents of step 1, after a series of CPU instructions in the series of POPs at the end, step 13, of the INT 21h dispatch in Figure 6-7. So this is how INT 21h function 52 returns SvVars to ENBX. If you want to see how DOS creates SvVars in the first place, you need to disassemble the DOS initialization code. You can have a peek from the appendix that lists an internal DOS function, INT 21h AX:1718, to get the caller's register structure, it returns a pointer to the structure in DS=SI. This word is the first subindex of a record at 110:84282. In fact, they are one and the same thing, the DOS call to save some things, a clear instruction pointer rather than through an INT 21h call case, so that in a table of INT 21h API 126 subroutines, 4282h does appear as the handler for subfunction 18h.



### A Very Brief Glance at File I/O

Next, let's look at a more interesting function. From Figure 6-8, the task for INT 21h AH=3F (Read File) is supposed to be located at EDC8:A839. The task for this function's task extension to examine in depth here, seems to pass on at the first two lines:

```

-> FDC8:A839
FDC8:A839 BCFD71      MOV     SI,71FD      ; offset of internal Read func
FDC8:A83C EB2DFE      CALL    A66C         ; see below

```

You know, but need to 3Fh expects a file handle. If you know to be written that the handles are associated with the current PSP. Examining the subroutine called at A66C it shows how DOS uses the passed in file handle:

```

-> FDC8:A66C
FDC8:A66C 2E8E06E730  MOV     ES,ES (3DE7) ; get DOS DS
FDC8:A671 268E063003  MOV     ES,ES (0350) ; ES < current PSP
FDC8:A676 263B1E3200  CMP     BX,ES (0032) ; PSP[32h] > # max open files
FDC8:A67B 7204        JB      A683
FDC8:A67D B006        MOV     AL,06        ; 6 invalid handle error
FDC8:A67F F9          BTC     ; set carry flag
FDC8:A680 C3          RET
FDC8:A683 76C43E3400  JES    DI,ES (0034) ; PSP[34h] > file handle tbl
FDC8:A686 03FB      ADD     DI,BX        ; use file handle as offset
FDC8:A68B C3          RET

```

In other words, DOS sometimes uses the current PSP to convert the passed-in file handle into a pointer to the kernel's internal File Table (see Chapter 8). Determining this pointer is a key entry into the system's kernel. Following the call, the DOS kernel function can examine the entry of the caller's workstation file. With some work (for example, in Raw mode) or most probably call down to a file table (see Chapter 8), which is with a normal file, a device driver must handle the call. Obviously, Reads do not necessarily get here in the first place, among them, being passed off to a disk access, as SMARTdrv5.1 file access, is the whole point of this task.

The subtask for DOS A66C is similar to the handle to the internal DOS file table. INT 21h AX=1200 (Get File Info) can be seen as a good example. A good example of the internal DOS functions call use a non-positive INT 21h AX=2100 to get pointer to a file's entry in the system. And the MOV DS:SI=32h task is a case seen in many times, such as of the AX=21h AX=1203 (Get DOS Data Segment) and the We keep a record of these INT 21h AX=1203 subfunctions, it's time to take a closer look.

### Tracing a DOS INT 2Fh Call

To examine the task for INT 21h AH=2Fh, we're going to reassemble the DOS INT 2Fh handler just as we did for INT 21h Read that we first used DE BUG to trace through a simple INT 2Fh call some core boot's DOS (N=2) in a debugger. We could do the same thing again for a simplified such as INT 21h AX=1200 (DOS kernel services installation) check, see appendix. But is there any way to automate some DE BUG's file. Can you perhaps trace through interrupts and locate a routine, interrupt chain without DE BUG's help?

### How Does DEBUG Trace Through an INT?

First you have to understand a little of how the DE BUG trace command works. The task is a set of *Undocumented DOS* and *Windows* computer by Tim Paterson (a debugging with extensive source code examples on the accompanying disk "UNDOC CHAP7" ASM). This is an excellent place to turn for a generic discussion of how DE BUG works.

The two excellent debuggers such as DEBUG and VMDB uses the single step feature built into the x86 system processors. When the processor's trace flag (TF) is enabled, the processor will execute only one instruction at a time. A debugger can install an INTR single step handler to receive the CPU's interrupt signal on every instruction.

The single step handler returns code 00h and leaving it enabled on entry to the single step mode would cause an endless loop. For this reason, the processor temporarily disables the trace flag (TF) several times per interrupt and re-enables tracing when the interrupt handler returns. In fact, the processor disables single step for *all* interrupts.

This is why any debugger must trace through an INTR to trace through an INTR. A debugger must do this by saving the breakpoint at the first instruction of the interrupt handler and then reenable single step mode. See *VMDB* and *VMDB* (see Crawford and Gosinger, *Disassembling the MS-DOS 3.31*) for more information. The MSD family of debuggers (started with the first edition of *Undocumented DOS*) happen to trace through INTR instructions.

## INTCHAIN

We can incorporate this knowledge into a program that single steps through an interrupt handler. INTRCHAIN is our own custom debugger. An INTR single step handler turns on the trace flag, then interrupts the processor, and then re-enables the processor's interrupt flag and turns off the trace flag. Because INTRCHAIN uses the same handler for an INTR, the processor calls the single step handler for all interrupts from the same interrupt handler. The handler saves away CS:IP whenever CS is changed, to make sure that the interrupt handler is changing to the previous handler. When the interrupt handler returns, INTRCHAIN is turned off, the trace flag INTRCHAIN prints out the interrupt chain saved by the single step handler.

The single step trace the same as in Figure 6-3 that SMARTDRV does back to anything of the DOS file *Reset* (hex code INTR 21h AH=00h). This is plainly visible in an INTRCHAIN trace of a call to this function, shown in Figure 6-10a.

**Figure 6-10a: INTRCHAIN Display for INT 21h AH=00h (Disk Reset)**

```
C:\WINDOWS\CHAP6>intrchain 21/0000
1387 instructions
Skipped over 4 INT

0B94 4200  HXDIER
07FA 15FA  SMARTDRV
CB41 0B29
0255 0025  J
474 058E  DB.55F55
F416 109F  DOS
F0CB 40FB  HRA
0070 06C5  IO
F0CB 8655  HRA
9070 C7D0  IO
F17 0763  HRA
C62C 0625  DB.55F55
07FA 1631  SMARTDRV
```

Notice that the interrupt passed by MSDOS SYS IO SYS and DBLSSYS, the call winds up back to SMARTDRV.

For a closer comparison with the DEBUG trace in Figure 6-3, Figure 6-10b presents sample INTRCHAIN output for a single tracing of INT 21h AH=62h call.

**Figure 6-10b: INTRCHAIN Display for INT 21h AH=62h**

```
C:\WINDOWS\CHAP6>intrchain 21/6200
77 instructions
```

```

DF93:32B6  WSCDEX
D7F9:15FA  SMARTDRv
C801:0800
D255:0023  0:
CC2E:058F  DBLSSYS
D116:109E  DOS
FDCB:40F8  WRA

```

Sure enough, this matches the interrupt chain we so laboriously traced back in Figure 6-3. INTCHAIN uses MAPC from Listing 6-2 to try to match up CSIP addresses with the names of resident ISRs and drivers. The addresses displayed by INTCHAIN can be passed to SYMDEF or DEBUGs for cross-referencing; this is the whole point of the program.

INTCHAIN can also trace through an XMS function or an arbitrary segment offset pointer. Actually, the program has little to do with interrupt chains as such. Rather than generate an actual INT instruction, you then have to deal with setting a breakpoint; the program just turns an INT XXh into a far call, and PUSH, to the handler for INT XXh. Thus, INTCHAIN won't trace through any INT generated inside the handler, such as the INT 21h call made with the INT 21h dispatch at Figure 6-7; this is generally what you want anyway.

### Listing 6-5: INTCHAIN.C

```

/*
INTCHAIN.C
Andrew Schulman, May 1993
Copyright (C) 1993 Andrew Schulman. All rights reserved.
*/

#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <dos.h>

typedef unsigned char BYTE;
typedef unsigned short WORD;
typedef unsigned long QWORD;

#ifdef __cplusplus
typedef void interrupt (far *INTRFUNC)(. );
#else
typedef void (interrupt far *INTRFUNC)(void);
#endif

typedef void (far *FARFUNC)(void);

#ifdef MK_FP
#define MK_FP(s,o) (((DWORD) s) << 16) + (o)
#endif

#define MK_INT(fp) (((DWORD) FP_SEG(fp)) << 4) + FP_OFF(fp)

#pragma pack(1)

typedef struct {
#ifdef __TURBOC__
WORD bp,di,si,ds,es,ds,cs,bx,ax;
#else

```

```

WORD es,ds,di,si,bp,sp,bx,dx,cs,ax; /* same as PUSHA */
#endif
WORD ip,cs,flags,
} REG_PARAMS;

#define INT_INSTR 0xCB
#define TRACE_FLAG 0x100

extern char *find_ומר(DWORD lin_addr); // in wop.c

void fail(const char *s) { puts(s); exit(1); }

#define MAX_ADDR 512

static WORD volatile instr = 0, int_instr = 0;
static WORD prev_seg = 0, my_seg = 0;
static void far *addr,
static int num_addr = 0;

void interrupt far single_step(REG_PARAMS r) // INT 1 handler
{
WORD seg,
BYTE far *fp,

if ((seg = r.cs) == my_seg) // ignore my own code
return,

fp = (BYTE far *) RW_FP(r.cs, r.ip),
if (fp[0] == INT_INSTR) // count INTs
int_instr++,
instr++,

if seg != prev_seg // if segment changed,
{ // assume we've chained
if (num_addr < MAX_ADDR)
addr[num_addr++] = (void far *) fp,
prev_seg = seg;
}
}

#define GET_FLAGS(reg) _asm ( pushf ), _asm ( pop reg )
#define SET_FLAGS(reg) _asm ( push reg ), _asm ( popf )

void set_flag(unsigned mask)
{
GET_FLAGS(ax);
_asm or ax, word ptr mask
SET_FLAGS(ax),
}

void clear_flag(unsigned mask)
{
GET_FLAGS(ax);
_asm mov bx, word ptr mask
_asm not bx
_asm and ax, bx
SET_FLAGS(ax),
}

FARPROC get_sms(void)
{
_asm mov ax, 4300h
_asm int 2fh
_asm cmp al, 80h
_asm je present

```

```

absent.
    fail("XMS not present");
present
    _asm mov ax, 4310h
    _asm int 21h
    _asm mov ax, bx
    _asm mov dx, es
    // return in DX AX
}

main(int argc, char *argv[])
{
    static int intrfunc = 0; /* make sure not in a register */
    INTRFUNC old_intr,
    FARFUNC func = (FARFUNC) 0;
    FARFUNC xms_func = (FARFUNC) 0;
    void far *fp,
    char *s,
    WORD intno, _ax, _bx, _cx, _dx;
    int a20off = 0,
    int i;

    puts("INTCHAIN 3.0 == Walks interrupt chains");
    puts("From undocumented DOS, 2nd edition (Addison-Wesley, 1993)");
    puts("Copyright (C) 1993 Andrew Schwab. All rights reserved.");

    if (argc < 2)
        fail("usage: intchain [ a20off ] [ intno, seg ofs/seg:bx/cdx ],");

    if (strcmp(strupr(argv[1]), "--A20OFF") == 0)
    {
        xms_func = get_xms(),
        a20off++,
        argv++;
    }

    // Figure out what code they want to generate:
    // an XMS call
    if (strcmp(strupr(argv[1]), "XMS", 5) == 0)
    {
        func = get_xms();
        sscanf(argv[2], "XMS %04x %04x %04x %04x",
            &_ax, &_bx, &_cx, &_dx);
        printf("Tracing XMS at %f\n", func);
    }
    // ... or a far (segment offset) call
    else if (strcmp(argv[1], "") == 0)
    {
        WORD seg, i;
        sscanf(argv[2], "%04x %04x %04x %04x %04x",
            &seg, &i, &_bx, &_cx, &_dx);
        func = (FARFUNC) 0;
        fp = seg, ofs = i;
        printf("Tracing function at %f\n", func);
    }
    // or an INT instruction
    else
    {
        sscanf(argv[2], "%02x %04x %04x %04x %04x",
            &intno, &_ax, &_bx, &_cx, &_dx);

        /* Single step, shouldn't go through INT, so turn the INT into
         * a PUSHF and far CALL */
        if (! (func = (FARFUNC) _dos_getvect(intno)))
            fail("INT unused");
        intrfunc++; // so do PUSHF when call func
    }
}

```

```

    printf("Tracing INT 20h AX=0x%x\n", intno, _ax);
}
/* (! (addr = (void far *) calloc(MAX_ADDR, sizeof(void far *)))
   fail("insufficient memory");
fp = (void far *) main;
my_seg = FP_SEG(fp);
old_sstep = _dos_getvect(1);
_dos_setvect(1, (INTRFUNC) single_step);
if (a20off)
{
    _asm mov ah, 6
    (^ams_func)(); // local disable A20 line
}
set_flag	TRACE_FLAG);
/* cail the code */
_asm mov ax, _ax
_asm mov bx, _bx
_asm mov cx, _cx
_asm mov dx, _dx
if (!intrfunc)
    _asm pushf
(!func());

clear f_eq(TRACE_FLAG);
_dos_setvect(1, old_sstep);
printf("In instructions\n", instr);
if (!int_instr)
    printf("Skipped over In INT\n", int_instr);
printf("\n");
for (i=0; i<num_addr; i++)
{
    s = find_owner(RK_LIN(addr[i]));
    printf("IPvt2s\n", addr[i], s? s : "");
}
if (num_addr == MAX_ADDR)
    fail("Overflow: very long INT chain");
return 0;
}

```

### Examining The INT 20h Chain

You can now use INTCHAIN to trace through a call to INT 20h AX=1200h, without using DEBUG. Figure 6-11 shows sample results. Note that the configuration is somewhat different from the one used to produce the INTCHAIN output for INT 20h AH=62h in Figure 6-10b.

**Figure 6-11: INTCHAIN Display for INT 20h AX=1200h**

```

C:\UNDOC2\CHAP6>intchain 20/1200
174 instructions
Skipped over 1 INT
1248 0007 NLSFUNC
109A 0980 PRINT
0116 0943 SHARE
0878 0285 DOSKEY
0894 3080 MSDISK
07FA 1568 SMARTDRV
0726-019F COMMAND
0725 0135 COMMAND
C801:0696
0725 018D COMMAND
FFFF:0FDB NMA

```

```

CB01:061E      J      DBLSSYSB
0255:0020      J      DBLSSYSB
CC2C:25E0      J      DBLSSYSB
0255:0028      J      DBLSSYSB
CC2C:0116      J      DBLSSYSB
D116:10C6      DOS
FDCB:64BD      NHA

```

Forward the goal of disassembling DOS, the essential piece of information here is the very last line as this gives the address 11D8:44BD of `MINDOS\SYSTEM\INT2FH` handler. We welcome back to this in a few moments.

The most noticeable feature of Figure 6-11 is the very long interrupt chain `INTERRUPT_PRINT_SHARE_IOCTL_KEY_MOUSE_SHARE_IOCTL_COMMAND_COM` +1 DoubleSpace, all take a `CALL` to process by the call. Possibly, even what is as you I see, an absolute entry to a `INT2FH AX=1200h` call compares the even `ISR` and device driver dumped out on `INT2FH` respect the call to see if it interests them. `INT2FH` chains can be extracted from `obj`, they are just calls, but then any interrupt handler as well as `CALL` such as the wrappers from Chapter 2 are involved. As noted earlier, Ralf Brown has suggested an alternate `INT2FH` protocol in an attempt to shorten long chains of handlers, so `CALL` for `INT2FH` calls appear.

Naturally, you can guess in the addresses displayed by `INTCHAIN` to a debugger since `DEBUG` or `SYMBOL`. For example, take the `0B94:308D` handler for `INT2FH` which `INTCHAIN` shows belong to the Microsoft CD-ROM Extensions.

```

-> 0B94:308D
0B94:308D 9C          PUSHF
0B94:308E 80FC15     CMP      AH,15
0B94:3091 7503     JNZ     3096
0B94:3093 EB6B     JMP     3100
0B94:3095 90          NOP
0B94:3096 80FC15     CMP      AH,15
0B94:3099 7503     JNZ     309E
0B94:309B EB09     JMP     30A6
0B94:309D 90          NOP
0B94:309E 80FC05     CMP      AH,05
< ...

```

You can see `MSDOEX` just checking for calls to `INT2FH AH=15h`. This makes perfect sense as `INT2FH AH=15h` is the network redirector protocol, and `MSDOEX` is a network redirector (see Chapter 8). `MSDOEX` checks for calls to `INT2FH AH=15h` to see again makes sense since this is the domain of `MSDOEX`. Act. see 30-31, *Dynamic MS-DOS Extension*. What about `INT2FH AH=05h`? As explained in the app. Ex this was an advanced user interface that also is used in programs now, kernel users interface, or to explain critical error or other errors seriously. External DOS programs such as `COMMAND.COM` issue `INT2FH AH=05h` calls, network redirectors such as `MSDOEX` handle calls to `INT2FH` with the value of `AX` with strings to display, such as `SCDRERR`. Not ready reading, I see `CALL` to `INT2FH` recording of `Handle's Mail`.

How about `INT2FH` for Windows? Figure 6-12 shows `INTCHAIN` output for the same configuration as Figure 6-11 except that `INTCHAIN` is running in a DOS box under Windows Enhanced mode.

**Figure 6-12: INTCHAIN Display for INT 2Fh AX=1200h Under Windows Enhanced Mode**

```

Tracing INT 2F AX=1200
175 instructions

1404:02A7      win
1264:0b68      WINICE
1248:00D7      MSLFUNC
109A:0980      PRINT

```

```

0F16 09-3  SHARE
0B78 0285  DOSKEY
0B94 3080  MSCDEX
07FA 1368  SMARTDRV
0726 019F  COMMAND
0725 0135  COMMAND
CB01 0696
1580 0045  win386
0725 019F  COMMAND
7F1F 0FD8  NRA
CB01 061E
0255 0070  J:
CC7C 25E0  DB,55FSB
0255 0028  J:
CC7C 0176  DB,55FSB
0F16 10C6  DOS
FDCB 4480  NRA

```

Not only are WIN.COM and WINFILE (the Soft ICE Windows debugger) added themselves to the command line INTELIVE, but notice that WIN386 has associated itself into the middle of the command line, covering itself with the help of a service interrupts from DOS boxes. Windows Enhanced mode is a mode designed for users of code that never show up in INTELIVE, at least in its present form. Many instructions, such as SHI and CHI, cause a jump into the Windows Virtual Machine Manager (VMM) to set up protected mode. This is possible because to a real mode DOS program like INTELIVE, a particular Windows enhanced mode looks like INTELIVE using the protected mode interrupt description table. A more sophisticated version of INTELIVE would need to be written to deal with Windows Enhanced mode. The same goes for INTELIVE Listing 6.13, which does at least one of the VMM instructions that Windows uses as a V86 breakpoint.

### The MSDOS.SYS and IO.SYS INT 2Fh Handlers

You can obtain the information on what which is the address in Figures 6.11 and 6.12. In the next section, we see the low memory segment for INT 2Fh when DOS INIT. In this configuration, the MSDOS.SYS INT 2Fh handler is located at 11B844BD; this code is shown with comments in Figure 6.13.

**Figure 6.13 MSDOS.SYS INT 2Fh Handler from MS-DOS 6.0**

```

<u>no fdcB 44bd
FD,8 4480 FB
FD,8 448E 80FC17  STI
FDCB 44C1 750A  CMP AH,11 ; 2F/17 network redirector call?
; no
; unupported functions /now here. Some external program like SHARE,
; MSLFBIH, or a redirector is supposed to handle these. If we got here,
; the external program must not be loaded, so it's an error. except
; if the user is doing a 2F/100 install, check in which case DOS
; will not return AX unchanged to indicate the software isn't installed.
FDCB 44C5 7403  OP AL,AL ; 2F/100 install check?
FDCB 44C5 7403  JZ 44CA ; yes, unsupported func, AX unchanged
FDCB 44C7 EBDCFF  CALL 44A6 ; no -- set carry flag for error
FDCB 44CA C0200  RETI 0002 ; sort of RETI without changing flags
FDCB 44CB 80FC10  CMP AH,10 ; 2F/10 SHARE call?
FDCB 44D0 74F8  JZ 44C3 ; yes: error
FDCB 44D2 80FC14  CMP AH,14 ; 2F/14 MSLFBIH call?
FDCB 44D5 74E8  JZ 44C3 ; yes: error
FDCB 44D7 80FC12  CMP AH,12 ; 2F/12 DOS internal function?
FDCB 44DA 7503  JNZ 44D0 ; no keep checking
FDCB 44DB E91102  JMP 44D0 ; yes goto fig 6-15a
FDCB 44DF 80FC16  CMP AH,16 ; 2F/16 Windows call or broadcast?

```





**Figure 6 15b: More MSDOS.SYS Code for INT 2Fh AH=12h**

```

-u fdc8 504a
fDc8 504A 55          PUSH  BP
fDc8 504B 8BEC       MOV   BP,SP
fDc8 504C 58        PUSH  BX
fDc8 504E 8B5E06     MOV   BX,[BP+06]      ; address of subfunc table
fDc8 5051 2E8A7F     MOV   BL,CS [BX]     ; number of valid subfuncs
fDc8 5054 3B5E04     Cmp   [BP+04],BL     ; caller's subfunction number
fDc8 5057 7317       JMB   5070           ; if too high, error
fDc8 5059 8A5E04     MOV   BL,[BP+04]    ; get subfunction
fDc8 505C 32FF     XOR   BH,BH
fDc8 505E 01E3     SHL  BX,1           ; turn into word offset
fDc8 5060 48       INC   BX            ; skip past # subfunctions
fDc8 5061 035E06     ADD  BX,[BP+06]     ; add in address of table
fDc8 5064 2E8B7F     MOV  BX,CS [BX]    ; pull out func ptr
fDc8 5066 895E06     MOV  [BP+06],BX    ; push on stack, RET to it
fDc8 506A 5B       POP   BX
fDc8 506B 5D       POP   BP
fDc8 506C 83C404     ADD  SP,+04
fDc8 506F 73       RET                                ; call subfunc via RET
fDc8 5070 5B       POP   BX                ; invalid subfunc come here
fDc8 5071 5D       POP   BP
fDc8 5072 C20600     RET   0006

```

This code is the ending table code in Figure 6 15b's set specifically related to INT 2Fh AH=12h calls. It is similar to the subfunction use in some subroutines. For example, the handler for INT 2Fh AH=0C calls has some subroutines. The top of Figure 6 15a shows that calling this subroutine requires you to set up function pointers on the stack including AX, which holds the function and subfunction. The table codes, for example, 1200h, hold the address of a table of function pointers. Hexadecimal 5051 holds the number of valid subfunctions; the rest of the table is an array of near function pointers to the appropriate handlers for each subfunction.

The subfunction in Figure 6 15b takes the caller's subfunction number, for example, the 06h in INT 2Fh, and compares it to the first byte of the table to see if it is within range. If yes, the code shifts the subfunction constant one word and adds to the address of the table; the value is incremented by a constant of the caller's test byte. The subroutine then pushes the function pointer out of the table, pushes the function pointer on the stack, and returns. (ret)

### Locating the INT 2Fh AH=12h Dispatch Table

This was somewhat difficult to do, but for our purposes the key piece of information is simply the location of the table, as this holds pointers to every INT 2Fh AH=12h subfunction. At the top of Figure 6 15a, there is a comment indicating that, in this configuration, the table is at 11D8 347C. The first byte of this value is the number of subfunctions. This is followed immediately by an array of 300 words holding function pointers to the various INT 2Fh AH=12h subfunctions.

```

db fdc8 3ff6 5ff6
fDc8 3f70          30          0
dw fdc8 3ff2
fDc8 3f70 470F 6E2E 4CBE 470B 9066 54E8 9342 9BEA
fDc8 3f80 6F2F 949F 038F 6B6A 6B53 4BCE 5030 9BE3
fDc8 3f90 5030 41F9 5011 9011 9927 9A76 A6A3 AB12
fDc8 3fA0 4282 A8BB AEC4 A97B A852 496C 4FD7 A9FC
;

```

Let's see if this is really the INT 2Fh AH=12h dispatch table. Earlier, it was noted that the subfunction 4282h that DOS uses to find out calling system's system code for INT 2Fh AX=1218h. Get C al to see if this is the case. Using SYMDEF to dump the table entry #186 confirms that this is correct:

```

-du fdc8 3ff2+(18*2)
fDc8 3fA0 4282 ....

```

The FTAB program (our Listing 6-4) can produce a nicer display of this same table. In fact, FTAB has an option to display tables such as this that keep the number of subfunctions as the first byte. The two commands shown in Listing 6-16 are thus equivalent. So that you have a handy 21-12 crib sheet to refer to, the entire table is shown, along with comments indicating the purpose of each subfunction.

**Figure 6-16: INT 2Fh AH=12h Dispatch Table Displayed by FTAB**

```
C:\VUNDO2\CHAP6>ftab fcb:3f7e 30 int2f12
C:\VUNDO2\CHAP6>ftab fcb 3f7d 7 int2f12
FCB:470E int2f12_00 ; install check
FCB:6E2E int2f12_01 ; close current file
FCB:4CBE int2f12_02 ; get interrupt addr
FCB:4708 int2f12_03 ; get dos data seg
FCB:9066 int2f12_04 ; normalize path separator
FCB:54F8 int2f12_05 ; output char
FCB:9342 int2f12_06 ; invoke crit err
FCB:985A int2f12_07 ; make disk buff most recently used
FCB:6F2F int2f12_08 ; decrement sft ref count
FCB:9A9F int2f12_09 ; flush and free disk buff
FCB:838F int2f12_0A ; perform crit err interrupt
FCB:686A int2f12_0B ; signal share violation
FCB:6853 int2f12_0C ; set fcb file's owner
FCB:48CE int2f12_0D ; get date and time
FCB:5030 int2f12_0E ; mark all disk buffer unreferenced
FCB:98E3 int2f12_0F ; make buffer most recently used
FCB:5030 int2f12_10 ; find unreferenced disk buffer
FCB:4FF9 int2f12_11 ; normalize asciz filename
FCB:5011 int2f12_12 ; strlen
FCB:9011 int2f12_13 ; toupper
FCB:9927 int2f12_14 ; _fstrcat
FCB:9A76 int2f12_15 ; flush buffer
FCB:A6A3 int2f12_16 ; get address of SFT entry
FCB:AB12 int2f12_17 ; set working drive
FCB:4282 int2f12_18 ; get caller's registers
FCB:A48B int2f12_19 ; set drive
FCB:AED9 int2f12_1A ; get file's drive
FCB:4978 int2f12_1B ; set year, length of february
FCB:4A12 int2f12_1C ; checksum memory
FCB:496C int2f12_1D ; sum memory
FCB:4FD7 int2f12_1E ; compare filenames
FCB:A9FC int2f12_1F ; build CDS
FCB:A66C int2f12_20 ; get JFI entry
FCB:AF48 int2f12_21 ; lruename
FCB:4434 int2f12_22 ; set extended err info
FCB:8147 int2f12_23 ; check if char dev
FCB:5030 int2f12_24 ; delay
FCB:501F int2f12_25 ; strlen
FCB:50B4 int2f12_26 ; open file
FCB:A72A int2f12_27 ; close file
FCB:50DA int2f12_28 ; move file pointer (lseek)
FCB:A839 int2f12_29 ; read file
FCB:5094 int2f12_2A ; set fastopen entry point
FCB:5117 int2f12_2B ; ioctl
FCB:5106 int2f12_2C ; get dev chain
FCB:5134 int2f12_2D ; get extended err code
FCB:5139 int2f12_2E ; get/set error table addresses
FCB:440D int2f12_2F ; nop
```

The work we've done to locate a INT 2Fh AH=12h was that we expected that many of the error function calls that JFS make internally would show up here. Indeed you can now see clearly that the C-ALL-4282 error has come from popped up in these experiments is actually INT 2Fh AH=12h. So...

Get is produced, variant CALL 4004 is actually INT 2Fh AX=1220h. Get J-E Entry D08h interlocks to ensure correct use of the functions in Figure 6-16, but, as already noted, it does so using a near CALL instruction. In INT D08h, DOS provides the INT form mostly for use by redirectors; see Chapter 8. Having this stack frame on the INT 2Fh AX=1217h functions definitely makes it much easier to understand what is doing, as the INT 2Fh functions are in which you are interested.

Recall that in Figures 6-11 and 6-12, the process of locating this table started by having the INT 2Fh AX=1217h call the INT 2Fh AX=1900h. This function, the DOS internal services install check, is used to determine if you are running with M-ETH to indicate that the services are present. The table indicates that the 0x4004 entry is a killer for this function. Let's unassemble it @ this address to check that the table makes sense.

```
> fdc8:470e
fdc8:470e 80FF          MOV     AL,FF
fdc8:4710 C3           RET
```

How is it? INT 2Fh AX=2000h, which is supposed to return with the DOS data segment in DS?

```
> fdc8:4708
fdc8:4708 2EBE1EE730   MOV     DS,CS [30E7]
fdc8:470D C3           RET
```

Let's do some work with the stack to look at a more interesting function. According to the appendix, INT 2Fh AX=1217h is a DOS working drive; the caller must push a zero-based drive number on the stack before calling this function. According to Figure 6-16, this function is located at FDC8:4702. Figure 6-17 shows comments I SIMMER assembled of this code.

#### Figure 6-17: INT 2Fh AX=1217h (Set Working Drive) in MS-DOS 6.0

```
> fdc8:4612
17: 88 points to DOS DB.
    mov     SysVar+18h at DOS:0026, 80 DOS:0047 is SysVar+21h
fdc8:4617 563A064700   CMP     AL,S:[0047] , SysVar+20h = LASTDRIVE
fdc8:4617 7202       JB     AB18 , %s drive < LASTDRIVE?
fdc8:4619 F9         STC     ; not: set carry flag, fail
fdc8:461A C3         RET
fdc8:461B 54       PUSH   BX , yes
fdc8:461C 50       PUSH   AX
fdc8:461D 563A163C00  DS     SI,S:[003C] , SysVar+16h = DOS ptr
fdc8:4622 8158     MOV     BL,%B , %Bh size of DOS entry
fdc8:4624 F6E3     MUL    BL
fdb8:4626 3F     ADD     SI,AX , DS SI = ptr to drive's COS
17: Here, 80A at DOS:0120, so DOS:05A2 is 80A+282h
fdc8:4628 568916A205  MOV     SS,[05A2],SI , move drive's COS ptr into
fdb8:462D 56871E4405  MOV     SS,[05A4],05 , DOS 80A+282h
fdc8:4632 5B     POP    AX
fdc8:4633 5B     POP    BX
fdc8:4634 FB     CLC
fdc8:4635 C3     RET
```

Let's start this out with the basic number on the stack; you may wonder how the code starts with a drive number in AX. Looking back at Figure 6-15, note that the generic INT 2Fh AX=1217h instruction is used on the stack. If you are aware of the caller's CNIP and flags, and know that the AX=1217h case is a function that does not expect a parameter on the stack, AX holds a good guess. This is not wise, but DOS makes an INT 2Fh AX=1217h call; this is just a start-up activity to complete DOS's internal call to the code for INT 2Fh AX=1217h, and that any parameter on the stack (INT 2Fh's version) would appear on the stack; see the appendix, actually appears in AX.

Even though the function involves fairly straightforward manipulation of DOS internal structures, the caller checks the drive number against the internal value of LASTDRIVE in SysVar. If

the drive number is valid, the function uses it as an index into the CDS array, a pointer to which is also contained in `SVSW`. The function then moves a pointer to the CDS entry into a DOS global variable. Changing this variable is basically what means to set DOS's working drive.

It is useful to see how DOS internally uses the `ASIDRIVE` variable in `SVSW` the CDS and other undocumented DOS features. Documentation of undocumented DOS are often, as far as the edition of this book is concerned, not given any consideration of DOS internals. For the CDS, `SVSW`, List of Lists, and other structures are not provided for our convenience. Like the hidden "junk screens" that software hobbyists use to do various secret screen finding. In fact, the CDS and so on are not so much undocumented DOS features as *external* DOS features that happen to be externally accessible through an undocumented interface. That undocumented DOS is often discussed without the surrounding context of DOS internals tends to obscure the real purpose of these structures.

For example, even though this chapter has often referred to the location of variables such as `CURR_PSP` as `SDA+100`, `ERR_AK`, `FILE_AK` as `SDA+1`, the `SDA` more really is more such thing as the `Swappable Data Area`. The `SDA` is merely an externally visible interface that Microsoft added rather late, on top of the DOS core segment (see "Origins of the SDA" in Chapter 8). I believe the `INT_21h` `ALL21` functions are just an undocumented external interface provided on top of some internal DOS functions for the convenience mostly of network redirectors. The external form of all form of these functions is the true one.

What has been accomplished here, basically, by locating the `INT_21h` `ALL21` dispatch table, we now acquire the means for *runtime* internal DOS functions. One of the nice things of the `INT_21h` dispatch table gives us means for *runtime* different locations in DOS. Rather than keep picking at disassembly of individual functions here and there, we can now turn around and do a full-blown disassembly of this entire code segment.

## Really Disassembling DOS

Everything we've done so far in DOS is in the same code segment, which in the particular configuration happens to be `1308h`. Of course, there are other parts of DOS on this segment, such as a good place to start. However, you disassemble the entire code segment at once, but still keep track of where the individual functions are located. For example, to a memory disassembler, it may not `FILE_AK` you would like to know where `Set PSP` function is located, where `SVSW` is found, et al. and so on.

You can use `DIRU` or `SVMDI` to produce a disassembly of this DOS code segment and use the `FILE` program to produce a new disassembly, the location of key functions within the segment. To merge the `FILE` output with the disassembly and, while we're about it, clean up and improve the disassembly, you can use `merge` as follows: a program named `SVMDI` written in `AWK`, a Unix pattern matching language, that is used here for text-processing tasks like this:

To examine the entire DOS code segment, you first need to know where to tell `DIRU` to start and stop processing. You can make a preliminary start by using the proper unassembly range by taking the first 32 segments of the `INT_21h` dispatch table (Figure 6-8) and the `INT_21h` `ALL21` dispatch table (Table 6-16), combining them, and sorting them by address:

```
C:\WINDOWS\CHAP6>type tmp.bat
@echo off
ftab fdc8:3e9e 6d INT21 > int21f.tmp
ftab fdc8:317d 30 INT2F_12 >> int21f.tmp
sort < int21f.tmp > int21f.log
```

```
C:\WINDOWS\CHAP6>tmp
```

```
C:\WINDOWS\CHAP6>type int21f.log
FCB 4052 INT21 33
```

```

FDCB 40A9 INT21_50
FDCB 40B5 INT21_51
FDCB 40B5 INT21_67
FDCB 40C1 INT21_6A
FDCB 42B2 INT2F_12_1B
FDCB 4400 INT21_1B
    etc
FDCB BDE9 INT21_5A
FDCB B1B5 INT21_61
FDCB B2D8 INT21_5C
FDCB B38F INT2F_12_0A

```

You can test the usefulness of INT21F.DIF if it's clear that you want DEBUC or SYMDIB to work. For example, at FDCB 40B2 and loading somewhere a bit after FDCB B38F, B3D0 is probably a jump instruction. You'll probably need to adjust the address range later on for DEBUC or SYMDIB, since it's not a jump. You can synthesize an assembly command into a batch script file, feed it to the debugger, and then use the debugger's output to a file:

```

C:\WINDOWS\CHAP6>type int21fscr
1. FDCB 40B2 B3D0
4

```

```

C:\WINDOWS\CHAP6>debug < int21fscr > int21f out

```

The `SYMDIB` debugger can DEBUC, produces its own SYMDIB puts segment override, and produces its own output on a separate line like DEBUC, but you must use the `SYMDIB /X` option and the `SYMDIB /SYMDIB` command prompt, which you wouldn't see if you redirected output to a file:

```

C:\WINDOWS\CHAP6>syndeb /x < int21fscr > int21f out

```

The `SYMDIB` command is fast. The `INT21F.DIF` showed be about 870k bytes, which should be only a small fraction of the size of the disk containing it. For `SYMDIB`, it's about 1/4 the size. The `DEBUC` command starts. One of the things `NTFDBG` can do is merge the `INT21F.DIF` file produced by DEBUC or SYMDIB with the `INT21F.DIF` file produced using `INT21F`.

## Windows Patches MS-DOS

**Actually, there's one interesting thing** you can do with the raw unassembly output from DEBUC or SYMDIB. Run the DEBUC unassembly script once under MS-DOS, then start Windows Enhanced mode, and rerun the DEBUC script again from inside a DOS box. Redirect DEBUC's output to a different file. This sequence gives you an easy way to examine the patches that Windows applies to MS-DOS. Just compare the two files, using `diff` or a similar utility. Any differences in this DOS code segment are the result of Windows patches.

```

C:\WINDOWS\CHAP6>debug < int21fscr > int21f out

```

```

C:\WINDOWS\CHAP6>win

```

```

::: from inside DOS box:

```

```

C:\WINDOWS\CHAP6>debug < int21fscr > int21f win

```

```

C:\WINDOWS\CHAP6>diff int21f out int21f win > int21f dif

```

The list of Windows patches in `INT21F.DIF` is incomplete, because it shows only one DOS code segment. Still, it does provide some idea of what is going on:

```

;; original MS-DOS code in INT 21h dispatch (see Figure 6-7 above)
< FDCB:41CF 50          PUSH AX
< FDCB:41CF B4B2       MOV AX,B2
< FDCB:41D1 CD2A       INT 2A
< FDCB:41D3 5B          POP AX

;; patched by Windows, 15AD belongs to WIN386.EXE
> FDCB:41CE 9ADAD0AD15  CALL 15AD 00A
> FDCB:41D3 90          NOP

;; original DOS code in a frequently called internal Begin Crit OI function
< FDCB:514B B8D180     MOV AX,B001
< FDCB:514E CD2A       INT 2A

;; patched by Windows
> FDCB:514B 9A6300AD15  CALL 15AD:00A3

;; original DOS code in a frequently called internal End Crit OI function
< FDCB:516B B8D181     MOV AX,B101
< FDCB:516E CD2A       INT 2A

;; patched by Windows
> FDCB:516B 9A7900AD15  CALL 15AD:0079

; .. etc.

```

The DOSMGR VxD built into WIN386 EXE applies these patches. When Windows exits, DOSMGR, of course, backs its changes out, restoring the original DOS code. As you can see, these patches have to do with DOS critical sections. DOSMGR wants DOS to call into the Windows VMM Begin Critical Section and End Critical Section functions. It's important to note that DOSMGR scans for the INT 2Ah instructions to patch, rather than using hardwired addresses. Thus, these patches should at least theoretically work with a different vendor's DOS.

The same before and after technique can be used to find DOS patches applied by other programs, such as MSCDEX. Programs that patch DOS can only be safely unoaded by a MARK RELEASE type of program that knows enough about these patches to back them out.

## Using NICEDBG

To run NICEDBG, use the companion disk. Feed it output from DEBUC or SYMDI. Optionally, you can supply a symbol table file, or a data file, address patch, such as FEAR patches. You can also supply NICEDBG's optional file of data file address patch. For example,

```

debug < int212f.scf > int212f.out
ftab fdcB:3e9e fd INT21 > int212f.log
ftab fdcB:3f7d 3D INT2F_12 >> int212f.log
nicedbg int212f.out int212f.log int212f.dat > int212f.txt

```

NICEDBG can take any file you give it as the output from DEBUC or SYMDI. The program makes several passes over the DEBUC file, replacing calls and jumps to meaningless address values like 4F5F, so you can find jumps to some real address patched by the user, such as INT2F\_12. NICEDBG also creates some useful labels for any other addresses that are the target of calls or jumps. Each target address itself contains a RET or JMP. NICEDBG changes the label to reflect this, so that it also generates a list of cross-references to each location.

For example, a jump to 00000000 from DEBUC looks like this:

```

FDCB:5126 9C          PUSHF
FDCB:5127 36          SS:

```

```

FOCB 5128 803E0C0000  CMP BYTE PTR EDI0C3,00
FOCB 5129 740F          JZ 513E
FOCB 512F EB01          JMP 5132
FOCB 5131 CF          IRET
FOCB 5132 0E          PUSH CS
FOCB 5133 EBF0FF      CALL 5131
FOCB 5136 50          PUSH AX
FOCB 5137 8B0180      MOV AX,8001
FOCB 513A C92A      INT 2A
FOCB 513C 5B          POP AX
FOCB 513D C3          RET
FOCB 513E EB01          JMP 5141
FOCB 5140 CF          IRET
FOCB 5141 0E          PUSH CS
FOCB 5142 EBF0FF      CALL 5140
FOCB 5145 C3          RET

```

This is not very promising looking. But NID FIDBC can transform this raw disassembly listing into something much more readable and useful:

```

; xref: FOCB:4304 FOCB:438B FOCB:4D7A
; func_5126:
FOCB 5126 9E          PUSHF
FOCB 5127 36803E0C0000  CMP BYTE PTR SS:[000C],00
FOCB 5128 740F          JZ jmp_513E -> loc_5141
FOCB 512F EB01          JMP loc_5132

; xref: FOCB 5133
FOCB 5131 CF          IRET

; xref: FOCB 512F
; loc_5132:
FOCB 5132 0E          PUSH CS
FOCB 5133 EBF0FF      CALL ret_5131
FOCB 5136 50          PUSH AX
FOCB 5137 8B0180      MOV AX,8001
FOCB 513A C92A      INT 2A
FOCB 513C 5B          POP AX
FOCB 513D C3          RET

; xref: FOCB 512B
FOCB:513E EB01          JMP loc_5141

; xref: FOCB 5142
FOCB 5140 CF          IRET

; xref: jmp_513E
; loc_5141:
FOCB:5141 0E          PUSH CS
FOCB 5142 EBF0FF      CALL ret_5140
FOCB 5145 C3          RET

```

Here are some of the changes that NID FIDBC made at various offsets in the code:

- 5126: Generated a cross-reference string `xref:` and a `func_5126` label because this location is called from three other places elsewhere in the code. To generate such labels and `xref:`, NID FIDBC's `fixup` must make multiple passes over the DEBUG listing.
- 5127: Cloned the `SS` operand together with the rest of the instruction at 5128.
- 5128: Replaced `JZ 513E` with `JZ jmp_513E -> loc_5141`. 513E is the target of the `JZ` at 5128. Because 513E itself does not unconditional `JMP`, NID FIDBC gives it a `jmp_513E` label.



Furthermore, NICEDBG chases down JMPs to JMPs and shows the final destination here (at 5141). The DOS code contains many JMPs to JMPs, knowing the ultimate destination of each JMP is to make the code a little easier to understand and also suggests a possible area for DOS optimization.

- 5131 Skipped a blank line because 5131 is the target of another instruction and thus starts a new block and because 5131 can not be reached from the preceding line at 512F, which does an unconditional JMP. NICEDBG generated a set-based offset of 01 to this location from 5133. Finally, because the code at this location does an immediate JRE, NICEDBG gave it a net 5131 label.
- 5132 Again, NICEDBG skipped a blank line because 5132 is target of another instruction and because 5132 can not be reached from the preceding line, which does an JRE. NICEDBG generated a low 5132 label.

NICEDBG uses `to` to specify targets of jumps, `from` to specify targets of CALLs, `loop` to specify targets of LOOPs, `ret` to specify code that ended its return via either RET or JRE, and `jump` to specify code that does an unconditional JMP. If the user supplies a symbol table (see `cd` name address pairs) associated with FILE, NICEDBG will use this as a source of labels.

NICEDBG AWK Listing 6-6 is the source code for this postprocessor for output from DEBUG or SYMDIB.

## What is AWK?

Since the reader is likely to be unfamiliar with AWK, a brief explanation of Listing 6-6 is probably called for. AWK reads in each line of text in one or more files and splits the line into fields. You can change the delimiters that AWK uses to decide where fields start and end, but it defaults to using white space, which is exactly what we need here. The fields are available to the program as `$1`, `$2`, and so on, up to `$NF` (`NF` is a built-in AWK variable that holds the number of fields). `$0` is the original line. For example, the line `"DC8 440D INT21 D"` is `$0 = DC8 440D`, `$1 = D`, and `INT21 D` is `$2` (and `$NF`).

Note too that AWK handles regular expressions (as also found in utilities such as `grep`). For example, the regular expression `(DES|S)` matches `"ES"`, `"DS"`, `"ES"`, or `"SS"` and `"\[*\]"` matches anything within square brackets. AWK also has associative arrays (just built-in hash tables, really), that can be indexed with strings (for example, `array["string"]`) as well as numbers. The presence of an item in an associative array can be tested with the `in` operator (for example, `fd` string in array).

The standard reference is *The AWK Programming Language* by Alfred Aho, Brian Kernighan, and Peter Weinberger, from the last letters of whose last names the language got its name. The high-level pattern-matching and array features of AWK make it possible to implement NICEDBG in about 200 lines of code.

NICEDBG.EXE on the accompanying disk was produced with the excellent AWK compiler from Thompson Automation. You can run the program without having AWK or understanding anything about it, but to modify the program, you would need Thompson AWK or another AWK interpreter or compiler. The popular MKS Toolkit comes with AWK, and many BBSs carry MAWK, a freely available, fast AWK interpreter by Mike Brennan.

NICEDBG processes each line in the DEBUG file. For example, consider the following line from a DEBUG listing:

```
DC8:5120 74DF          JZ     513E
```

AWK reads this line into fields skinned by spaces. The *nth* field is referred to as *\$n*.

```

$1 $0CB.5120      Address of the instruction
$2 740F          Instruction opcode bytes -
$3 JZ           Instruction operator
$4 513E         Instruction operand

```

Of course, not every instruction looks quite like this. For example,

```

$1 $2          $3 $4 $5 $6
F0CB:5126 9C   PUSHF
F0CB:5127 56   SS:
F0CB 5128 B05E0C0000 CMP BYTE PTR [000C],00

```

in any case, NCFDBG's AWK can rely on \$1 as the address of the instruction and on \$3 as either the instruction operator or, when using DBRUCs rather than SYMDEFs, something like a segment override.

Before processing the DBRUC file, NCFDBG reads in the optional symbol table and data files. NCFDBG uses IN12121.F00C or any similarly formatted file to build a table of names, called *ftab*, corresponding to symbols. (Obviously, the program runs through each line in IN12121.001, or any file assembly listing produced by DBRUCs or SYMDEFs, to see if the line's segment (other address) is in the table.)

NCFDBG makes three passes over the DBRUC file.

**Pass 1.** NCFDBG looks for any calls, jumps, or loops in the code and adds the target of the call, jump, or loop to *ftab*, which it will later use to generate labels. Simplifying considerably, the AWK code looks like this:

```

if ($3 ~ /CALL/) ftab[$4] = "func_" $4;
if ($3 ~ /LOOP/) ftab[$4] = "loop_" $4;
if ($3 ~ /J.*/) ftab[$4] = "loc_" $4;

```

in pass 1, NCFDBG also constructs the *jmptrab* for resolving JMPs to JMPs:

```

if ($5 ~ /JMP/) jmptrab[$1] $4, # jmptrab[SOURCE] + TARGET

```

**Pass 2.** The second time through the DBRUC file, NCFDBG builds its *ret* table and also improves some of the labels generated in pass 1. A label such as *jmp\_XXXX* or *ret\_XXXX* indicating that location XXXX does an unconditional JMP or RET is generally more useful than a label such as *loc\_XXXX*, indicating that XXXX is the target of a jump. Thus, if pass 1 assigned a location a name, and that location does a JMP or a RET, NCFDBG changes *ftab* to reflect this:

```

if ($5 ~ /*RET/) $6 ($1 in ftab) ftab[$1] = ret_ $1;
if ($5 ~ /JMP/) $6 ($1 in ftab) ftab[$1] = jmp_ $1;

```

Also in pass 2, NCFDBG looks for code that may be "not reached" — that is, not accessible from any other location in the listing, of course, the code might be called from outside the disassembly range. If the previous assembly did an unconditional JMP or RET, and if there are no labels at the current address, *ftab*[\$1] computes, indicating that \$1 is not the target of a jump, call, or loop. NCFDBG adds \$1 to a *not\_reached* array:

```

if (!did_jumpet == 1) $6 (" $1 in ftab) not_reached[$1]++;
did_jumpet = 0;
if ($3 ~ /*RET|JMP/) did_jumpet = 1;

```

**Pass 3.** In its final pass over the DBRUC listing, NCFDBG prints out the new, improved listing:

- If the current line was not reached, *not\_reached* = \$1 or *not\_reached*. NCFDBG indicates this with a comment: this might be a sign either of data, or of "dead code."

- If the current line has one or more labels (e.g., STARTUP.DRV), NICEDBG prints them out one per line. Either the user will have supplied these labels in a symbol table file, or NICEDBG will have generated them automatically in passes 1 and 2.
- If any other line of code jumped, called, or looped to this line of code, NICEDBG displays the cross reference.
- If the code itself calls, jumps, or loops to some other line of code, NICEDBG replaces the numeric target with a name from table 54 (table 54).
- If the target of the call, jump, or loop is a line of code that itself does an unconditional JMP, NICEDBG chases down any JMP's to JMP's (see the resolve\_jump\_jmp() function in Listing 6-6).
- If the user supplied a list of data address name pairs, NICEDBG uses it to replace addresses such as 0330<sub>h</sub> with names provided by the user, such as CURRENT\_PSP. Recall that an earlier inspection of GetPSP and SetPSP (Figures 6-3 and 6-4) showed that those functions do nothing more than get and set a word at offset 0330<sub>h</sub> in DOS.DOS. Thus, a occurrence of 0330<sub>h</sub> in an assembly of DOS can probably be replaced with a name such as CURRENT\_PSP. Likewise, examination of INT\_21h API 55h (Figure 6-8) showed that DOS.0009<sub>h</sub> is STARTUP.DRV, and that 0337<sub>h</sub> is BRK\_FLAG. You can feed this sort of information to NICEDBG in a file of data address name pairs, such as IN1212.DAT.

```
0069      STARTUP.DRV
0330      CURR_PSP
0337      BRK_FLAG
30E7      DOS_DS
1030      IN_WINSE
033E      MACHINE_ID
0321      IN_DOS
0584      USER_SP
0586      USER_SS
0320      CRIT_ERR
1211      DOS_HIGH
```

But note that NICEDBG's replacements of, for example, [0330] with CURRENT\_PSP are very simplistic. The program merely does a blind global search and replace. Thus you should be cautious about what you put in a NICE.DAT.DAT file.

- If the debugger that SYMDEF was used to produce NICEDBG's output, NICEDBG saves away a byte sequence (searched on the current line, CS: EIP) and uses the AWK substitution construct to smack it into its proper place on the next line.

### Listing 6-6: NICEDBG.AWK

```
# NICEDBG.AWK -- Produces nicer output from DEBUG input and symbol table
# usage: nicedbg symbol dbgfile logfile
# example: nicedbg int212f.log int212f.out int212f.lst

# get offset from seg:ofs
function get_off(addr) { split(addr, so, ":"), return so[2]; }

function mk_fp(ofs) { return seg "-" ofs; } # make seg:ofs farptr

function get_ftab_name(addr) { # get name from table
  if (addr ~ SEG_OFS)
    addr = mk_fp(addr); # table indexed by seg:ofs
  if (! (addr in ftab))
    return addr; # not there -- return unchanged
  split(ftab[addr], label, ",");
  return label[1]; # just return first name w/ 1
}
```

```

function resolve_jmp_jmp(src) { # JMP to JMP to . .
  if (! (src in jmptab))
    return;
  if (done[src])
    return done[src];
  # If get here, haven't seen this one yet
  target = target2 = jmptab[src];
  while (target in jmptab) {
    target2 = jmptab[target];
    if (target2 == target) # endless loop
      break;
    if (target2 == src) # cycle
      break;
    if (target2 in done) { # we've seen this part already
      target2 = done[target2];
      break;
    }
    target = target2;
  }
  done[src] = target2;
  return target2;
}

function hexix { return 0 + ! Dx - s); } # relies on Thompson AMR

BEGIN {
  print N[DEBUG] "Makes nicer output from DEBUG input and symbol table",
  print "From s_jndocumented DOS", 2nd edition (Addison-Wesley, 1993);
  print "Copyright (c) 1995 Andrew Schulman All rights reserved \n";
  if (ARGC < 2) {
    print usage nice dbgfile {syntab} {datfile} {outfile}
    print "example nice dbg int2121 out int2121 log int2121 lai"
    did_anything = 0;
    exit;
  }
  else did_anything = 1;

  # commonly-used regular expressions
  SQ_BRACK [ [ '\]' ], # anything within square brackets
  SEG_OFS = /^[^:]/, # has a : in it
  SEG_OVERRIDE /^[DES]$/, # CS or DS or ES, or SS;
  CALL_OR_JMP = /CALL|LOOP|J*/, # CALL, LOOP, JMP, J*

  # read in optional symbol-table file
  # lines in syntab file look like: xxxx:yyyy name
  if (ARGC < 2) {
    while (getline < ARGV[2]) # for each line in symbol table
      tab[$1] = tab[$1] $2 " ", # put name into table for seg:ofs
    close(ARGV[2]);
  }

  # read in optional data file
  # lines in data file look like: xxxx name
  # example: 0521 IN_DOS
  if (ARGC < 3) {
    while (getline < ARGV[3])
      data[$1] = $2,
    close(ARGV[3]);
  }

  ARGC = 2; # finished with sym, dat file
  dbgfile = ARGV[1]; # switch over to DEBUG file

  # debug file looks like. xxxx:yyyy XXXXX op operands
  # example: FDCB:4052 3C06 CMP AL,06 ; comments

```

```

while (getline < dbgfile) { # make pass 1 through debug file
  if ($1 ~ SEG_OFS) {
    split($1, so, "-");
    if (" seg" {
      seg = so[1]; # get segment for later use
      start = hex(so[2]);
    }
    else
      stop = so[2], # take last one
    }
  if ($3 ~ CALL_OR_JMP) {
    if ($4 ~ /\[^\[\*\]\]/FAR/) # don't do [xxxx] or xxxx.yyyy etc
      continue;
    # should also ignore e.g. CALL $I
    if ($3 ~ /JMP/)
      jmtab[get_off($1)] = $4, # jmtab for resolving JMP JRP
    if (" mk_fp($4) in ftab) # put call/jmp target into table
      ftab[mk_fp($4)] = (($3 ~ /CALL/) ? "func_" :
        ($3 ~ /LOOP/) ? "loop_" : "loc_") $4,
    }
}
close(dbgfile);
stop = hex(stop);

# pass 2: build cross-ref table, improve some label names, etc.
while (getline < dbgfile) {
  if (did_jmpret == 1) && (" $1 in ftab))
    not_reached[$1]++; # prev line did JMP/RET, but no label, so
  did_jmpret = 0, # "not reached", may be data or dead code

  if ($3 ~ /JMP|JMP/) {
    did_jmpret = 1,
    if (" $1 in ftab) # if target is a ret/jmp, change label name
      ftab[$1] = (($3 ~ /JMP/) ? jmp_" : ret_) get_off($1),
    # oops, this will also replace labels supplied in sym file!
  }
  # below *not* "else if" - JMP handled both places
  # build xref table and outside-range table
  if (($3 ~ CALL_OR_JMP) && ($4 ~ SQ_BRACK) && ($5 !~ SQ_BRACK)) {
    if ($6 ~ /FAR/)
      outside[$5]++;
    else if ($4 ~ SEG_OFS)
      outside[$4]++;
    else {
      off = hex($4),
      if ((off start) || (off stop))
        outside[off]++;
    }
  }
  if ($6 !~ /\[^\[\*\]\]/FAR/) # don't do [xxxx] or xxxx.yyyy
    xref[mk_fp($4)] = xref[mk_fp($4)] get_ftab_name($1) " ",
  }
}
close(dbgfile);

# pass 3: for each line in dbg file
while (" $1 SEG_OFS) {
  print, getline;
  if (" $D) exit,
}

japline = ,

# indicate if this is possible unreachable (dead) code; show
# cross-reference (xref) table; show all labels for this address

```

```

if ($1 in not_reached) {                                # possible dead code
  print "
  print ";;; not reached?";
}
else if ($1 in ftab) {                                  # if segment:offset in table
  print
  if (xref[$1])
    print "xref: " xref[$1] # show xref
  nf = split(ftab[$1], label, ",");
  for (i=1; i<nf;i++)
    if (label[i])
      printf("%24s:\n", " ", label[i]), # show all labels for this addr
  ftab_found[$1] = 1;
}

# if a CALL, LOOP, or some kind of JMP, show eventual destination
# of any JMP, JNP, and possibly replace number address with string name
if ($3 ~ (CALL|DR|JMP)) {
  if ($4 ~ (FAR)) {
    if ($4 in jmptab)
      jmpline = " " get_ftab_name(resolve_jmp_jmp($4));
    $4 = get_ftab_name($4); # replace number with name
  }
}

# cheap replacement of [xxxx] with names from data file
if (match($0, $0 BRACK)) # match sets RSTART, RLENGTH
  if (faddr = substr($0, RSTART+1, RLENGTH-2) in data)
    sub $0_BRACK, data[faddr], $0, # sub() does substitution

# get rid of DEBUG segment override ugliness
if ($3 ~ SEG_OVERRIDE) {
  override_addr = $1; # save to use on next line
  byte = $2;
  override = $3;
}
else if (override_addr) {
  $1 = override_addr; override_addr = "";
  $2 = byte $2;
  sub(/\[?, override \[, $0, # plug in override
}

# print out (possibly altered) line
if (! override_addr) {
  printf("%s\t%-15s\t", $1, $2);
  for (i = 3; i < NF; i++)
    printf "%s\t", $i);
  if (! jmpline)
    printf "\t", jmpline);
  printf("\n");
}
}

# print list of CALL, JMP, etc references outside disasm range
END {
  if (did_anything) {
    printf("\n, outside range %s %04x-%04x\n", seg, start, stop);
    for (x in outside)
      printf("%s = (%x - SEG_OFS) ? \"%s\" : \"%04x\" \"\n", x);
    # should suppress following if within a not_reached block?
    printf("\n, possible unresolved labels:\n");
    for (x in ftab)
      if (! (x in ftab_found))
        printf("%s\n", ftab[x]);
  }
}
}

```

With output from `DEBUG` in `INT21F.DAT`, a symbol table produced by `FILAB` in `INT21F.LIB`, and the optional data file `INT21F.DAT`, you can produce a nice-looking disassembly of the entire `MSDOS.SYS` code segment `INT21F.FSI` with

```
nicedbg int21f out int21f log int21f dat > int21f.lst
```

We will examine this `INT21F.FSI` file in more detail momentarily, but the following excerpt provides some idea of what `NICEDBG` produces.

```

INT2F_12_1B:
FDCB:4282 2E8E1EE73D    MOV DS,CS DOS_DS
FDCB:4287 C5368405    LDS SI,USER_3P
FDCB:428B C3          RET
; ...
INT2F_34:
FDCB:4059 EB26F5    CALL INT2F_12_1B
FDCB:405C C744022103    MOV WORD PTR [SI+02],1H_DOS
FDCB:4061 8C541D    MOV [SI+10],55
FDCB:4064 C3          RET

```

This is quite usable. You can see that `INT 21h AH=34h` Get InDOS Flag Address calls the code for `INT 21h AX=121Bh`, Get Callers Registers, and then moves `DOS_DS:1B` into the caller's `SI` register. This is just as you would expect.

You could make it a little more readable by going into `INT21F.LIB` and taking the only partially useful names, such as `INT2F_34` and `INT2F_12_1B` produced by `FILAB`, and replacing them with more evocative names, such as `GET_INDOS_34` and `GET_STACKPTR_121B`. But this is left as an exercise for the reader, who may or may not know a little DOS inner workings by heart and not require such a crutch. The point is simply that you can manually change or add to `INT21F.LIB` as you like, via `SYMBOL` functions. For example, you can add the following two functions that you already know about from running `INICMAIN`.

```

FDCB:40FB INT2F_DISPATCH
FDCB:448D INT2F_DISPATCH

```

Please note that `INT21F.FSI` is *not* included on the accompanying disk, as redistributing a large piece of `MSDOS` would obviously violate Microsoft's copyright. However, it should be easy for readers to produce their own personal copies, given the instructions in this chapter. To quickly summarize the steps involved in producing `INT21F.FSI`:

1. `INICMAIN 21_6200` and use last line to locate `DOS:INT 21h` handler.
2. `DEBUG` or `SYMBOL` to disassemble `INT 21h` handler, locate dispatch table.
3. `Read FILE = INT 21h dispatch table > tmpfile`.
4. `INICMAIN 21_7900` and use last line to locate `DOS:INT 21h` handler.
5. `DEBUG` or `SYMBOL` to disassemble `INT 21h` handler, locate dispatch table.
6. `Read FILE = INT 21h dispatch table > tmpfile`.
7. `SCRIP < tmpfile > symfile`.
8. `Topic < path > at option of symbol to create script for DEBUG or SYMBOL`.
9. `DEBUG < script > outfile`.
10. Optionally create data file.
11. Optionally change and add to `symfile`.
12. `NICEDBG outfile symfile [datafile] > lstfile`.
13. Check `outsid range comment` at end of batch. Possibly alter script, and goto step 9.

The last point needs a bit of explanation. Because code and data are interpreted within `DOS`, `DEBUG`, and `SYMBOL` are likely to encounter data that they will misinterpret as code. This is avoid-

could also throw on the unassembly of valid code further on in memory. The result is that IN12.2F.D51 may contain, for example, several CALLs to func\_9024, but, instead of showing code at offset 9024, there's instead some bogus-looking instruction at offset 9023h. NICEDEBUG will list such "pov" or "void" instructions at the end of the listing; you can use this to split the DEBUG or SYM file into segments into two or more parts. For example, let's say that there are valid-looking calls to func\_9024, but to func\_9024 itself. If the original DEBUG script contained the following command:

```

u fdc8:4052 b500
; or, say, this, it would make DEBUG restart unassembly at offset 9024h
u fdc8:4052 9024
u fdc8:9024 b500
    
```

At this point, of course, you may find idea of postprocessing DEBUG output a little ridiculous. You'll have to look to commercial disassemblers such as V Communications' Sourcerer.

But even with the DOS disassembled just one MSDOS.SYS code segment. You can apply the same technique to other parts of MSDOS: the outside range list produced by NICEDEBUG's help() here), to DR.DOS or to NetWare's NBTX code.

### Examining a Few DOS Functions

Let's look at a small portion of the MS-DOS 6.0 disassembly produced by DEBUG with a little help from FLAR, INCELAN, and NICEDEBUG. Figure 6-18 shows the code for a few simple DOS functions.

**Figure 6-18: MS-DOS 6.0 Code for Functions 34h, 52h, 1Fh, 32h, and 00h**

```

                                INT2F_34:
FDC8:4059  EB26F5          CALL INT2F_12_18
FDC8:405A  C744022103     MOV Word Ptr [SI+02],0321
FDC8:4061  8C5410         MOV [SI+10],55
FDC8:4064  C3            RET

                                INT2F_52:
FDC8:4065  E81A75        CALL INT2F_12_18
FDC8:4068  C744022600     MOV Word Ptr [SI+02],0026
FDC8:406B  8C5410         MOV [SI+10],55
FDC8:4070  C3            RET

                                INT2F_1F:
FDC8:4071  B200          MOV BL,00

                                INT2F_32:
FDC8:4073  76           PUSH 55
FDC8:4074  7F           POP 05
FDC8:4075  8AC2         MOV AL,BL
FDC8:4077  E84150        CALL INT2F_12_19
FDC8:407A  7222         JB Loc_409E
FDC8:407C  C43EA205     LES DI,[05A2]
FDC8:4080  26F6454480   TEST Byte Ptr ES:[DI+44],80
FDC8:4085  7517         JNZ Loc_409E
FDC8:4087  E80001        CALL Func_515A
FDC8:408A  E83749        CALL Func_96C4
FDC8:408E  E8CADC        CALL Func_515A
FDC8:4090  720C         JB Loc_409E
FDC8:4092  EBEDF4        CALL INT2F_12_18
FDC8:4095  896C02        MOV [SI+02],BF
FDC8:4098  8C440E        MOV [SI+0E],E5
FDC8:409B  32C0         XOR AL,AL
FDC8:409D  C3            RET
    
```





```

FDCB 4080 26F6454480 TEST Byte Ptr ES:[01+44],80 , CDSE[43-44h] = flags
FDCB 4085 7517 JNZ loc_409E ; if net/redir drive, fail
FDCB 4087 E88003 CALL func_513A ; enter crit #1 (2A/8001)
FDCB 408A E83749 CALL func_96C4 ; ES:BP get DFB from CDSE[45h]
FDCB 408D E8CA03 CALL func_513A ; wait crit #1 (2A/8101)
FDCB 4090 720C JB loc_409E ; fail?
FDCB 4092 E8ED14 CALL INT2F_12_18 ; get caller's regs
FDCB 4095 896C02 MOV [S1+02],BP ; caller's BP
FDCB 4098 8C440E MOV [S1+0E],ES ; caller's DS
FDCB 409B 32CD XOR AL,AL ; al = 0 for success
FDCB 409D C3 RET

```

Internal function *examine* back in Figure 6-18 is INT 21h AH=0Dh Disk Rese. The function does its real work inside the call to turn 9A34 (see above) which occupies all buffers, calling the once-to-HDD buffer function INT 21h AX=215h. But note in Figure 6-18 that Disk Rese also calls INT 21h AX=1120h which is the network redirector Flush All Disk Buffers function. This provides a good illustration of how the network redirector works as a series of *hooks* in DOS. At various key moments, DOS issues an INT 21h AH=11h call; any installed redirector can pick up the call and do what it needs to do (see Chapter 8).

One of the findings that probably first came from the DOS code shown in this chapter, but which becomes clearer as you examine the INT 21h INT 13h code, is that hooks play an important role in DOS. In addition to the INT 21h AH=11h redirector interface, DOS also checks the SMART hooks. These, however, simply function as a total-differing manner from the redirector—see SMARTHOOK at Listing 8-27 in Chapter 8. Of course, many DOS functions get passed down to installable device drivers. In DOS code calls these, however, using the Strategic and Interrupt pointers in the device driver header (see Chapter 7).

Remember also that external programs probably hook many of these DOS calls. You saw earlier, for example, that SMARTDRV and DRISPACE hook the Disk Rese call. Thus it is a little misleading to say the INT 21h AH=0Dh handler or MIRROR.sys is *simple*. When examining the code for a DOS function, it is important to remember that DOS isn't just the code in MIRROR.sys and IO.sys. Games, for example, tend to intercept many of these calls with a DOS extension; you are likely to find one at some point. This not only means understanding the role of programs such as Windows, SMARTDRV, MSCDEX, DONKEY, and CHN2WAP, but also understanding where (in Microsoft programs) is *ms* Straker, NetWare, and 486MAN from. A good example of this, as we saw in Chapter 4, is how a fairly simple NetPnP driver suddenly takes on new meaning and complexity when Novell NetWare is running.

## Examining the DOS Lseek Function

As a more extensive, but still remarkably self-contained, example, let's examine the DOS Move File To End function (INT 21h AH=42h), frequently known as *lseek* after its C Unix equivalent. We had occasion to examine the DOS code for this function while working on Chapter 8 of this book. An example of the network redirector specification in Chapter 8, in discussing the redirector INT 21h AX=121h Seek From End function, asserted that DOS never calls this function. Since this was based merely on empirical evidence, we never saw 21\_1121 called, it made sense to examine the DOS code—seems that DOS did not contain a call to INT 21h AX=1121h.

Our surprise: the DOS code for *seek* did contain a call to this INT 21h function. It turns out that DOS *only* calls the redirector's Seek From End function under a special set of circumstances having to do with network FCBS and various SMART modes. Frankly, we still don't quite understand this. In any case, the rest of the code for INT 21h AH=42h is fairly straightforward, yet long enough to be a little more interesting than the feeble little examples we've seen so far. In addition, there is some interesting Windows-related code in DOS that we'll encounter along the way.

Before we examine the disassembly listing for INT 21h AH=42h, recall that the function has the following specification:

#### Move File Pointer

**Input:**  
 AH = 42h  
 AL = method (0 = from beginning, 1 = from current pos, 2 = from end)  
 BX = file handle  
 CX-DX = hi/lo offset from beginning, current, or end  
 INT 21h

**Output success:**  
 Carry clear  
 DX:AX = new hi/lo position

**Output failure:**  
 Carry set  
 AX = error value (5 = invalid function, 6 = invalid handle)

Microsoft's DOS programmer's reference further notes that

A program should never attempt to move the file pointer to a position before the start of the file. Although this action does not generate an error during the move, it does generate an error on a subsequent read or write operation. A program can move the file pointer beyond the end of the file. On a subsequent write operation, MS-DOS writes data to the given position in the file, filling the gap between the previous end of the file and the given position with undefined data. This is a common way to reserve file space without writing to the file.

This suggests that almost any CVDV parameters to seek are valid. Indeed, as we're about to see, the code slices that we're about to see use the CVDV parameter into the file's SEI entry. The hard part is getting the SEI entry to make sense of the code listing you'll need to know the following objects at the SEI. For further information, see the appendix under INT 21h AH=52h.

02h	WORD	open mode
05h	WORD	device info word
11h	DWORD	file size
13h	DWORD	current file position
2Fh	WORD	machine number (Windows VR ID)

Figure 6-20 shows the DOS code for INT 21h AH=42h, Move File Pointer. Many explanatory comments were added by hand to the code generated by NCUDDC.

#### Figure 6-20: MS-DOS 6.0 Code for INT 21h AH=42h (Iseek)

```
; xref: FDCB,50B5 FDCB,9B52 FDCB,9BC1 FDCB,9E9C
                                INT21_42
FDCB:AB45 EBE100          CALL func_A929, TURNS BX HANDLE INTO
                                ; ES:DI SFT (see fig. 6-21)
; xref: FDCB:AB84
                                loc_AB48
FDCB:AB48 7302          JNB loc_AB4C
FDCB:AB4A E89C          JMP jmp_A7EA -> loc_43E0 ; couldn't fail!
; xref: loc_AB48
                                loc_AB4C
FDCB:AB4C 3C02          CMP AL,02
                                ; ES:DI=valid SFT entry
FDCB:AB4E 760A          JBE loc_AB5A
                                ; which move method"
FDCB:AB50 36C606280501  MOV Byte Ptr SS:[0323],01, SDA+3=error locus
FDCB:AB56 B001          MOV AL,01
                                ; 1=invalid function

; note why jmp jmp in DOS code:
; AB58 -> A7EA -> A706 -> A7D4 -> A716 -> A6FB -> 43E0
; usually to use short jmp, but is it still worth it?"
```

```

; but can it ever be changed?
; xref jmp_A8A0
FDCB A858  E090      jmp_A858-
                    JMP jmp_A7EA  > loc_43ED  , fail'

; xref FDCB A84E
FDCB A85A  5C0F      loc_A85A:
FDCB A85C  720A      CMP AL,D1
FDCB A85E  777B      JB loc_A868          , below = 0
                    JA loc_A87B          ; above = 2

FDCB A860  26035515    , handling seek method #1: from current pos
FDCB A864  26154017    ADD BX,ES [D1+15]    , SFT->file_pos
                    ADC CX,ES [D1+17]
                    , fail: through to method #0

; xref FDCB A85C FDCB A88A
FDCB A868  B0C1      loc_A868
FDCB A86A  92        MOV AX,CX          , #0, from beginning
                    XCHG AX,BX          ; BX:AX <- CX:BX

; xref FDCB A8A9
FDCB A86B  26894515    inc_A86B
FDCB A86E  26895517    MOV ES [D1+15],AX   , update SFT->file_pos
FDCB A871  F8FF99    MOV ES [D1+17],BX
FDCB A874  895406    CALL INT2F_32 1B    , get caller's regs
FDCB A876  895406    MOV [SI+06],DX      , move into caller's DX
                    ; later on, inc_43FB does MOV [SI], AX
                    ; see table 0-2 for caller reg struct

; xref: jmp_A8EF
FDCB A879  1BA7      jmp_A879
                    JMP jmp_A877  -> loc_43E6  , success'

; xref FDCB A85E
FDCB A87B  26F6450680    loc_A87B:          ; #2 from end
FDCB A880  750A      TEST Byte Ptr ES:[D1+06],80 , dev info: NETWORK
                    JNZ loc_A88C

; xref FDCB A891 FDCB A8A2
FDCB A882  26035511    loc_A882
FDCB A885  26154015    ADD BX,ES [D1+11]   , SFT->file_size
FDCB A888  26154015    ADC CX,ES [D1+15]   , CX BX += file_size
FDCB A88A  EB0C      JMP loc_A868         , go to method #0

; xref FDCB A880
FDCB A88C  75DE      loc_A88C:          ; this is a network drive!
                    ; ; this is seek method #2 (from end of file), and network bit is set
                    ; ; in S11 DOS may call a network redirector's 2F1121 Seek From End
                    ; ; handler, but only if some strange conditions are met. It can't
                    ; ; be an FCB open, and certain SHARE bits must be set

FDCB A88E  26F6450380    TEST Byte Ptr ES [D1+03],80 , open mode: FCB'
FDCB A891  75EF      JNZ loc_A88E        ; an FCB open

                    ; ; this is not an FCB open ;;;
FDCB A893  768B4502    MOV AX,ES [D1+02]   , open mode
FDCB A897  25FC00    AND AX,00F0
FDCB A89A  304C00    CMP AX,0040         , OPEN_SHARE_DENYNONE
FDCB A89D  7405      JZ DO_2F_1121      , redir seek from end
FDCB A89F  5D3C00    CMP AX,0030         ; OPEN_SHARE_DENYREAD
FDCB A8A2  75DE      JNZ loc_A88E        , no update caller's regs

; xref: FDCB A89D

```

```

DO_2F_1121
FDCB:A8A4  8B2111      MOV AX,1121 ; Call network redirector's
FDCB:A8A7  C02F      INT 2F      ; Seek from End function
FDCB:A8A9  73CD      JNB loc_A86B ; update caller's DX:AX from SFT

; xref: jmp_A8F9
jmp_A8AB
FDCB:A8AB  EBAB      JMP jmp_A858 > loc_43ED ; fail

```

As you can see, the code verifies the caller passes an invalid file handle in BX or an invalid seek method in AL. But once in this function, essentially `dos111.c` can be updated to current position in the kernel's FAT table. Even for files on network drives, DOS attempts to seek directly in seek call without calling a network redirector for assistance. We can see that the previous operation in this way:

```

sft = get_sft(handle) // see below
if (seek from begin) then set sft->file_pos = new_pos
if (seek from end) then (signed) new_pos += file_size, goto seek from begin
if (seek from current) then new_pos += sft->file_pos, goto seek from begin
set caller's new_pos (DX:AX) = sft->file_pos

```

We haven't updated the very first line of the INT 21h API 42h handler, however, since DOS calls a subprogram to seek a file first. API 42h to turn the caller's file handle into an SFT entry in the FSDT table is shown in Figure 6-21. The code for the API 42h entry set to be very interesting, because it's a common MS-DOS interaction with Windows. As early as Windows 3.11, some old FAT file drives it is same as the function is also called by other parts of DOS, including the code for functions 3fh and 69h.

**Figure 6-21. MS-DOS 6.0 Code To Verify SFT Virtual Machine ID**

```

; xref: INT21_3E INT21_68 FDCB:A7E5 INT21_42 FDCB:AB01 FDCB:A007
; func_A929
FDCB:A929  8B1EFC      CALL func_A62A ; into ES:D1 SFT (fig. 6-22)
FDCB:A92C  721C      JB ret_A94A ; percolate error up
; we'd handle, but it could be for another DOS box!
FDCB:A92E  50      PUSH AX
FDCB:A92F  36F60630'001 TEST Byte Pir $5 IN_41H3E,01
FDCB:A935  7404      JZ loc_A93B
FDCB:A937  33CD      XOR AX,AX
FDCB:A939  EB0B      JMP loc_A943

; xref: FDCB:A935
loc_A93B
FDCB:A93B  56A13ED3   MOV AX,$$ MACHINE_ID ; Windows running
FDCB:A93F  2670452F   CMP AX,ES [D1+2F] ; SFT's share_machine

; xref: FDCB:A939
loc_A943:
FDCB:A943  5B      POP AX ; okay
FDCB:A944  7501      JNZ loc_A947
FDCB:A946  C3      RET

; xref: FDCB:A944
loc_A947:
FDCB:A947  B006      MOV AL,06 ; failure
FDCB:A949  19      SFC ; "invalid handle"

; xref: FDCB:A92C
ret_A94A:
FDCB:A94A  C3      RET

```

This code deals with the fact that under Windows Enhanced mode, it is possible to have multiple processes (called DOS boxes) that happen to have the same PSP ID (though note that SYSTEMINFO has a UniqueDOS\_PSP setting). Normally, the current PSP and a file handle are sufficient to specify a particular file. Under Windows Enhanced mode, the current virtual machine VM ID is also needed to specify an open file.

In this section, DOSBox checks whether Windows Enhanced mode is running (see Chapter 1), sets the DOS ID to match the INI WIN32 flag, gets the current VM ID (see Chapter 1) to see what a DOS MGR VxD patches DOS\_MACHINE\_ID word with the current VM ID, and compares the current VM ID against the machine ID held at offset 21h in the SEI. If the SEI's machine ID does not match the current VM ID, the current VM ID's entry fails with error code 0, as if the handle in BX were not a file system object pointer, but it belonged to another process that happened to have the same PSP in another DOS box.

We then take the file handle in BX into an SEI entry in INDI. First we convert the file's WinAPI file handle to 22, which then turns the BX handle (which is only a pointer to the current PSP's job table table) into a HLL pointer equivalent to INI 21h WinAPI. Then we use the HLL pointer in our SEI index, and then turns the SEI index into an SEI entry, which is then INI 21h WinAPI. The disassembly starts off with DOS's INI 21h WinAPI 1220h and then WinAPI's other address in the string.

**Figure 6-22. MS-DOS 6.0 Code To Turn File Handle into SFT Pointer**

```

; xref FDCB 4FD1 func A62A loc A671 loc A6EA loc_A7DD FDCB A90F FDCB:A924
; INT21_12_2D
FDCB A6D0 2F8FD6D730 MOV ES,CS DOS_DS ; get DOS_DS
FDCB A672 2681067003 MOV [5,ES CURR_PSP ; use current PSP
FDCB A677 265B1F5200 CMP BX,ES [DOS2] ; # files in JFT
FDCB:A67C 7204 JB loc_A672
FDCB A67E B0D6 MOV AL,D6 ; invalid handle

; xref FDCB A657
; loc_A62D
FDCB A62D 19 STI
FDCB A621 C3 RET

; xref FDCB A67C
; loc_A672
FDCB A627 76C43E3400 LES DI,ES [DOS4] ; file handle < # files
FDCB A627 057B ADD DI,BX ; JFT ptr in PSP
; add on BX handle

; xref FDCB:A62D
; ret_A629
FDCB A629 C3 RET ; return ptr -> SFT idx

; ; code to turn handle in BX into SFT entry in ES:DI ; ;
; xref FDCB 4EDC INI21_4400_03 INI21_4402_03 FDCB 67DD INI21_440A 1
; FDCB 757B func_A929 FDCB:B27B
; func_A62A
FDCB A624 EBE0FF CALL INI21_12_2D ; turn BX handle->ES:DI JFT
FDCB A62D 72FA JB ret_A629
FDCB A62F 26805DFF CMP Byte Ptr ES [DI],FF ; unused
FDCB A635 7504 JNZ loc_A639
FDCB A635 B0D6 MOV AL,D6 ; invalid handle
FDCB A637 EBE7 JMP loc_A62D ; fail

; xref FDCB A635
; loc_A639
FDCB A639 55 PUSH BX
FDCB:A63A 268A1D MOV BL,ES[DI] ; JFT entry -> SFT index
FDCB A63D 32FF XOR BH,BH

```

```

FDCB:A65F  E80200      CALL INT2F_12_16      ; SFT index -> SFT ES DI
FDCB:A662  5B           POP BX
FDCB:A663  C3           RET

; xref: FDCB:6DF1 FDCB:A516 FDCB:A63F FDCB:A686
; INT2F_12_16
FDCB:A644  2E8E06D73D   MOV ES,CS DOS_05     ; SFT ndx -> ES:DI SFT
FDCB:A649  26C43E2400   LES DI,ES [002A]     ; get DOS_05
; SysVars+4 > first SFT

; xref: FDCB:A65E
; .loc A64E
FDCB:A64E  263B5004     CMP BX,ES [DI+04]    ; walk SFT chain
FDCB:A652  720E        JB Loc_A662          ; SFT # files
; in this table
FDCB:A654  26205004     SUB BX,ES [DI+04]    ; subtract #files this SFT
FDCB:A658  26C43B      LES DI,ES [DI]      ; follow nkd list
FDCB:A65B  83FFFF      CMP DI,DI           ; end of Sfts?
FDCB:A65E  73EE        JNZ Loc_A64E        ; loop to next SFT
FDCB:A660  F9         STC                 ; invalid SFT index
FDCB:A661  C3         RET                 ; fail?

; xref: FDCB:A652
; .loc A662
FDCB:A662  50         PUSH AX             ; in this SFT
FDCB:A663  B83B00     MOV AX,003B        ; SFT each size entry
FDCB:A666  F6E3      MUL BL             ;
FDCB:A668  D3FB      ADD DI,AX          ; offset of this entry
FDCB:A66A  5B         POP AX             ;
FDCB:A66B  83C706     ADD DI,+06        ; skip past SFT header
FDCB:A66E  C3         RET

```

The basic section starts by BX handle. If it entry 2E 1220 SFT index SFT entry 2E 1216.

Recall that the file handle to BX is really a device to be written to PSP's FILE. Thus the code for INT 21h AX 1220h gets the current PSP pointer to a regular CURR PSP global DOS variable and checks PSW 0032h which holds the maximum number of file handles a valid entry has. If the handle to BX is the file handle maximum, i.e. the FILE size, then this code gets a pointer to the FILE from PSW 0034h. It takes BX entry FILE pointer, adding a far pointer to PSW 0034h, the file's FILE entry.

Each FILE entry is a single byte that holds an index into the SFT. If FILE is selected as a SFT entry. The code in Figure 6-22 ensures that if the caller hasn't passed a file handle to use, corresponding FILE entry is unused.

If DOS has a valid SFT index, it passes it to a function equivalent to INT 21h AX 1216h, which returns a pointer to a corresponding SFT entry. From the snippet above we can see how this code works. DOS gets a pointer to the first SFT from SysVars+4. It walks the SFT chain comparing the SFT index to the number of files each SFT entry holds to find the right one. DOS then multiplies the result by SFT index by 3B. The size of an SFT entry in this version of DOS, and adds it onto the start of this SFT, to form an SFT entry.

That's it. We've now examined the DOS code for back in its entirety. We've seen how the specification for INT 21h FILE entry is actually implemented in working code, how DOS gets from a file handle to a FILE entry, how it finds the FILE and how it extracts the SFT entry and sets the entry file pointer to a device to be written to. The Windows API documentation remembers that this DOS SFT is possibly the only "back" to the important third party extensions such as NetWare, how the seek function. Our discussion of the DOS kernel neglects to deal with whatever changes they might make to the behavior of seek.

We have now presented a detailed random selection of extremely simple DOS functions viewed in relation from key third party DOS extensions. To properly discuss this simple DBL C disassembly of 80 gigabytes of DOS code, it would require an entire book. Exact properties of the code, which

examining its interactions with resident software such as SmartDrv, Windows, and NetWare could easily be the subject of several books. For further in-depth discussions of this code, see Chappell's *DOS Internals* and Mike Podanofski's *Dissecting DOS: A Cable Level Look at the DOS Operating System*. This book's main focus is described in more detail later in this chapter.

### Other Parts of DOS

As noted earlier, NIDBDB places an "outside range" list at the end of a disassembly listing. This list indicates locations that are called or jumped to in the listing, but which don't themselves appear in the listing. This list provides additional addresses for disassembly by DEBUG or SIMDB.

To exempt the disassembly of the MSDOS.SYS code segment, the function IN12FH.DIB.VIC.H. As you know from the earlier investigation in Figure 6-13, the INT 2Fh handler in MSDOS.SYS jumps to the handler in IO.SYS. Here is how this shows up in the IN12FH.IN1 file produced by NIDBDB:

```

j ref  FDCB 440A FDCB 462F FDCB 4687 FDCB 46ED
                                jmp_440F
FDCB,440F  EA05007000          JMP 0070 0005

```

```

*
** outside range FDCB 4045-8800
** 0070 0005
*

```

You can use the address 0070 0005 as the starting point for a disassembly of the IO.SYS code:

```

C:\UNDOC2\CHAP6>symdeb
=> 0070 0005 0005
0070 0005 EA95087000      JMP 0070 0893
=> 0070 0893 0893
0070 0893 2E1F2E06      JMP FAR CS [06E6]
  dd 0070 06e6 06e6
0070 06e6 7FFF 1502
  w 7FFF 1502
FFFF 1502 80FC15      CMP AH,15
FFFF 1505 7613      JZ 151A
FFFF 1507 80FC08      CMP AH,08
FFFF 150A 7438      JZ 1547
FFFF 150C 80FC16      CMP AH,16
FFFF 150F 7479      JZ 158A
FFFF 1511 80FC4A      CMP AH,4A
FFFF 1514 7503      JNZ 1519
FFFF 1516 E9A700      JMP 15C0
FFFF 1519 CF      IRET
-q

```

```

C:\UNDOC2\CHAP6>type io scr
w 7FFF 1502 1519
q

```

```

C:\UNDOC2\CHAP6>symdeb /a < io,scr > io.out

```

```

E:\UNDOC2\CHAP6>nicedbg io out > io.lst

```

```

C:\UNDOC2\CHAP6>type io.lst
*
** outside range 7FFF:1502-1519
** 151A
** 1547
** 158A
** 15C0

```



Now, of course we expand the unassembly range for SYMDI.B based on the addresses in the outside range list. Also we cannot yet create a file with symbolic names:

```
C:\UNDOC2\CHAP6>type fo scr
u ffff:1302 13c0
■

C:\UNDOC2\CHAP6>type fo sym
ffff:1302 10_INT2F
ffff:131A 10_INT2F_15
ffff:1347 10_INT2F_08
ffff:138A 10_INT2F_16
ffff:13CD 10_INT2F_4A

C:\UNDOC2\CHAP6>symdeb /x < fo.scr > fo.out

C:\UNDOC2\CHAP6>nicedbg fo.out fo.sym > fo.lst

C:\UNDOC2\CHAP6>type fo.lst
?
?
;; outside range ffff:1302-13CD
;; 0070 0898
;; 174E
```

We continue to this way until no unsolved references remain. As noted earlier, sometimes DEBUG and SYMDI.B get their own track because of data residing in the middle of a code segment. Based on the NLIST.DW symbols of a file, it is common need to split a single command in a DEBUG script into two or more separate `u` commands.

(3) course, the addresses shown here for disassembly in memory of MS-DOS.SYS and IO.SYS also work for a lot of other resident software. In Figure 6-11, for example, we saw SMARTDRV, MSCDEX, HOOKKEY, SHARE, PRINT, COMMAND.COM and several things piped out on the INT 21h chain. You can also find a lot of the addresses displayed by NLIST.DW in DEBUG or SYMDI.B for disassembly and process by directing output with NLIST.B.

However, since each coder to disassemble separate programs such as SMARTDRV, MSCDEX, COMMAND.COM and PRINT, or look rather than in memory, because these programs don't move the segments in some configurations and the DOS kernel, PRINT, in particular, is probably the most finished piece of DOS, as the system, so a TSR writers learned their craft. You can also see a disassembler such as Sourcery to examine these programs.

Given the ability to reverse engineer DOS, an almost infinite amount of information on DOS programming is readily available. To answer some questions about DOS, look at the code running on your machine, and one may solve problems with a new approach. It is that what someone else's configuration may include to you, another. Applications such DOS, DOS changes, through so much information from one version to the next. Describing software based on its source code, whether supplied or disassembled, can be the only way to see the way things fit out what the software really does, or it can be dangerous, revealing features that may change. There are no certainties here. Your best bet is to examine the sources, even but recognize how it may change, either because of future versions or because of unforeseen interactions with other software.

## Am I Going to Jail for This?

Many programmers ask for bits about the legality of what we've been doing in this chapter. Programmers frequently think that disassembling Microsoft's code is legal, and even that it's something a firm lawyer would let a technician do for an offense punishable by a soft prison sentence. We had better look into this now.

The following discussion of the legalities of disassembly was not written by an attorney, and should not in any way be viewed as legal advice. However, I have benefited enormously from discussions with Gene K. Landy, a partner at the law firm of Shapiro, Israe, & Werner, P.C., in Boston. Any errors and misconceptions of course remain mine.

Landy is the author of a superb book/disk package, *The Software Developer's and Marketer's Legal Companion*, published by Addison-Wesley (1993), which includes several extremely useful discussions of reverse engineering. Chapter 1 discusses reverse engineering in the context of copyright—including the important *Sega v. Accolade* case. Chapter 2 discusses software trade secrets and confidentiality agreements. Chapter 11 covers shrink-wrap licenses and warranties and the standard shrink-wrap license limitation on reverse engineering, noting the important case of *Vault v. Quaid*. This is a fine book that every software developer will want to have in these troubled, legally complex times.

While some programmers believe that you can wind up behind bars just for having seen the CPU instructions of the signature of the INT 2H dispatch code. Quite simply because the standard license agreement that comes with all Microsoft products states as plain as day:

3. OTHER RESTRICTIONS You may not reverse engineer, decompile, or disassemble the software.

The rest of that license agreement states that "this is a legal agreement between you (either an individual or a company) and Microsoft Corporation. By opening the sealed software package you are agreeing to be bound by the terms of this Agreement."

Well, but without doubt, if I disassemble Microsoft software, you have entered into a binding agreement of disassembly, even if disassembly were otherwise a legitimate activity, right?

No. You may have some positional advantage with shrink-wrap licenses, but because of the secrecy of their use, the law does not see that you've decided issues of shrink-wrap licenses have spread further dissemination than they deserve. As Landy explains in his chapter on shrink-wrap licenses:

In contrast to a contract, a shrink-wrap license is a system of acceptance or rejection. If you accept the contract, you can open the envelope; if you reject it, you return the package. For a contract, but not this "contract," someone works. Does the law really allow the licensor to force the user to this choice?

At the least, when a contract, even from its origin with century roots to the present, is the mirror that awakens the senses and a "meeting of the minds." In a classic contract, two parties are bargained to, the law takes place as agreed. While the sale of goods in all states except Louisiana is now governed by a state statute, the Uniform Commercial Code, the same concept is carried over. A contract and its terms are agreed, set out, or at the time of the sale. The problem with the Shrink Wrap License is that the retail software sales, over and done with before the customer is presented with one-sided terms of the Shrink Wrap license. After the sale is already made, it is too late to try to impose adverse terms.

Similarly, Raymond J. Summer's excellent textbook, *The Law of Computer Technology* notes that "The attempt to limit the capabilities of the common purchaser by virtue of a printed form included within the product package is unlikely to be successful."

How about the specific shrink-wrap license limitation against disassembly and reverse engineering. Two important cases have held that shrink-wrap or tear-me-open license agreements cannot be used to

outlaw reverse engineering. Both Landy's book and Nimmer's discusses the important case of *Intel v. Quindt* 198, 1988. The state of Louisiana had enacted special legislation to stabilize various aspects of shrink-wrap licenses, including the restriction on reverse engineering. A California corporation took Quindt, a Canadian corporation, to court in Louisiana to test the advantage of this exceptional law. A stormy trial for Auld, but fortunately for those who think that disassembly is an important consumer right, the court ruled that the Louisiana statute was preempted by federal law. A similar Louisiana statute has been repealed.

No Microsoft's shrink-wrap license limitation against disassembly probably saw worth in paper it's printed on.

How about the law of trade secrets? To begin with, reverse engineering is already one of the few *legitimus viis* to discover a trade secret. The Uniform Trade Secrets Act (UTSA) adopted in the mid 1980s by almost all states says explicitly that discoverers through reverse engineering is a proper means of gaining access to non-patented trade secrets. Knowing one of the main books on intellectual property, more or less at random, we find Roger E. Schuchter, *Unfair Trade Practice and Intellectual Property*, pp. 135-136, states added:

#### REVERSE ENGINEERING IS NOT IMPROPER MEANS

Many products are manufactured pursuant to plans or with technologies that are trade secrets and thus sold to the public at large. In some cases the method and manufacture of these items may be discovered by careful study of the object. Typical methods of discovery include taking the product apart or performing experiments on it. This process of analysis is usually called "reverse engineering." *Anonymous case had that contra-ante merita is non-meritoproprietas of contra-ante-merita.* Risk of discovery by reverse engineering is a risk that a firm takes when it chooses to rely on trade secret protection for a valuable commercial asset. Note that if a firm secures patent protection for a new device or manufacturing process it is protected against reverse engineering. This is one of the most important differences between patent and trade secret protection.

Given that MS-DOS is not patented, in two patent numbers, 4,958,069 and 5,109,433, in the front of all Microsoft's manuals are in a form of data compression, as used, for example, in Microsoft's file compressors, and then it seems to be quite straightforward. As a trade secret, it is concerned, reverse-secreting, a critical way. The rationale here is that trade secret law is basically about the loyalty of employees or others who receive important business information in confidence. You violate trade secret law when you breach, including or exposing, a relationship of trust. One does not violate another's trust by disassembling a product purchased on the open market.

So far, the shrink-wrap license statement against disassembling seems effective, and trade secrets do not disassemble books. What about the fact that MS-DOS is copyrighted? Does copyright law permit us to study how DOS works internally, and then build products based on this new-found knowledge, for example, how to violate Microsoft's copyright to reproduce how DOS preloads DBLSPACE.BIN as MS-DOS 6.0 and then create a replacement for DBLSPACE.BIN that supports the same interface?

Disassembly is sometimes regarded as a form of copying, translation from one medium to another, or, in language to another, and therefore, as possible copyright infringement. However, disassembly for the purposes of achieving compatibility is generally regarded as "fair use." An important decision by the United States Court of Appeals in the Ninth Circuit in *Sega v. Accol*, August 1992, overruling a lower court's ruling, held that Accol's use of knowledge reverse engineered from the Sega Genesis system did not violate Sega's copyright and constituted fair use. According to this court (as quoted in *UNIX Review*, May 1993),

We could do that only if reverse disassembly is the only way to gain access to the ideas and functional elements embodied in a copyrighted computer program, and where there is a

legitimate reason for seeking such access: disassembly is a fair use of the copyrighted work, as a matter of law.

The importance of Segal's verdict was underlined in a comment in *Microprocessor Report* (December 9, 1992): "For the industry, it may be a breath of a deep sigh of relief. No longer are we awaiting copyright suits as we seek to understand the parameters of an undocumented 'Int 21 call'."

Nature's natural enemies of the industry brought a sigh of relief on bearing the appeals court's check. It is not by a group calling itself the Business Equipment Manufacturers, which includes IBM, Intel, and Microsoft, is seeking stronger protection against reverse engineering. Arguing for greater protection of the cause engineering is the so-called American Committee for Interoperable Systems, which includes IBM, Microsofts, Amfdahl and Gups & Technologies (see "Reverse Engineering Reversals," *Update*, May 1993).

It is essential for the purpose of achieving compatibility is okay, and this by the way is also the a. The only other method of the I.C.S. directive is software protection, then how about this book's question from freemish instances. Have we violated Microsoft's copyright by reprinting several chunks of code from MS-DOS and Windows in this book?

As you can see, programs of copyright computer programs are considered to be "library works." Without a copyright notice that a compiled program without its source code, merits being called a later work of the public. Library work means anything at all in the context of computer software, it is not possible to create a library. Our inclusion of brief excerpts from disassembly listings is a really a form of a usual quotation, which is one of the oldest forms of fair use (see William S. Strong, *The Copyright Book*, 4th edition, Chapter 8).

It is not the fact that the I.C.S. has received a DMCA 2 tool which Microsoft works with the copyright of MS-DOS. Microsoft has made an effort to secure MS-DOS against disassembly, particularly since DMCA's ability to trace into an INI 21h or INI 21h call.

## Use the Source, Luke!

Is there any alternative to disassembly? One alternative is, of course, to rely entirely on the vendor's documentation and not to wonder whether this documentation is an accurate reflection of the actual software. But as the reader has probably recognized by now, relying on vendor documentation has as many risks as does relying on reverse-engineered software that has been discovered through disassembly.

Depending on what you are interested in, there may be another, less-alternative to disassembly: source code.

For example, programmers in custom communities aren't really about how the operating system behaves in a certain circumstance, but about what the compiler (in this case, the library (R11) does. There is a particular confusion among many programmers about the difference between a "lib" or C and a DOS file name. For example, it is often called the DOS Set Handle Command (INI 21h AH=67) and then wonder why the C program function still have C attributes such as yet to be created up by a code, which is the R11 source code. Both Microsoft C and Borland C++ come with R11 source code.

Sometimes, rather than having specific questions about MS-DOS, programmers are just curious about how operating systems work in general. In this case, the best approach is probably to study one of the several excellent books available on the design and implementation of UNIX. Some of these books are *Design of the UNIX Operating System* and *Building the UNIX System Architecture* (previously titled *Pseudocode for UNIX*). Others, such as *Flaninbach's monolithic Operating System: Design and Implementation*, *MINIX*, and *Conix's Operating System Design: The UNIX Approach*, come with complete source code for UNIX workalikes. Despite the numerous differences between DOS and UNIX, these books should be required reading for anyone planning to delve into DOS internals.

DOS's handling of various processes, files, devices, and so on, can often best be understood by contrasting it with the design and implementation of a well-understood system such as UNIX.

For a more serious disassembly approach to operating system design and implementation, another alternative to disassembly of MS-DOS is to examine the source code of a suitable commercial DOS workalike. The Real DOS from Amiga, Inc. Software, Richmond, VA, has Steve Jones's superb documentation of DOS's internals. For an excellent discussion of making a fully runnable DOS, see Steve's article "DOS Meets Real DOS" in the February 1992 *Unpublished System Programmer*. General Software's Utility SDK and Device Driver SDK come with complete source code in C for versions of DOS with a CHKDISK, FORMAT, FDISK, DISKCOPY, and ROM-DOS 5 from DataLight. Artagon, WA, is also available with source code.

Last but not least, Mike Podolowski's *makeprowd* sections show when RADOS (an open-source DOS available with fully commented assembly language source code, Podolowski is currently writing a full-length book on RADOS, *Real DOS: A Code to Code to DOS Operating System*, which will be available in 1994. While obviously not identical to the MS-DOS source, this source code may be more than adequate for your needs. For example, Figure 6-23 shows the implementation of INT 21h services 50h, 51h, and 52h from RADOS ASM.

**Figure 6-23: RadOS Implementation of INT 21h AH=50h, 51h, and 52h**

```

.....
; 50h Set PSP Address
;-----
; bx contains PSP address to use
;-----
_GetPSPAddress
mov word ptr [ _RxDOS_CurrentPSP ], bx ; Seg Pointer to current PSP
ret

.....
; 51h Get PSP Address
;-----
; bx contains PSP address to use
;-----
_GetPSPAddress
mov bx, word ptr [ _RxDOS_CurrentPSP ] ; Seg Pointer of current PSP
RetCalleeStackFrame es, si
mov word ptr es:[ _BX ][ si ], bx
ret

.....
; 52h Get Dos Data Table Pointer
;-----
; es:bx returns pointer to dos device parameter block
;-----
; DOS Undocumented Feature
;-----
_GetDosDataTablePtr
RetCalleeStackFrame es, si
mov word ptr es:[ _ExtraSegment ][ si ], ds
mov word ptr es:[ _BX ][ si ], offset _RxDOS_pDPB
clc
ret

```

There are no magic pixes here, really, how else could Get and Set PSP be implemented, any way that we can see? This is a courtesy to the MS-DOS, and that have to this code in the chapter might have saved us a lot of trouble.

More interestingly, Figure 6-24 shows the RADOS implementation of back the MS-DOS implementation of `write` (`write`). Figure 6-20. The RADOS code provides a useful guide to MS-DOS disassembly.

Figure 6-24: RxDOS implementation of INT 21h AH=42h (lseek)

```

.....
; 42h Lseek (Move) File Pointer
;-----
; al      move method
; bx      handle
; cx:dx   distance to move pointer
;-----
MoveFilePointer:
Entry
def _method, ax
def _handle, bx
ddef _moveDistance, cx, dx
ddef _newPosition

mov ax, bx ; handle
call MapAppToSysHandles ; map to internal handle info
call FindSFTByHandle ; get corresponding SFT (esi: di)
jc _moveFilePointer_36 ; if could not find ->

getdarg cx, dx, _moveDistance
mov ax, word ptr [ _method ][ bp ]
Goto SEEK_BEG, _moveFilePointer_beg
Goto SEEK_CUR, _moveFilePointer_cur
Goto SEEK_END, _moveFilePointer_end
SetError -1, _moveFilePointer_36

; seek from end
;-----
_moveFilePointer_end
add dx, word ptr es [ sttFileSize_low ][ di ]
adc cx, word ptr es [ sttFileSize_high ][ di ]
jmp short _moveFilePointer_beg

; seek from current position
;-----
_moveFilePointer_cur
add dx, word ptr es [ sttFilePosition_low ][ di ]
adc cx, word ptr es [ sttFilePosition_high ][ di ]
; jmp short _moveFilePointer_beg

; seek from beg on on
;-----
_moveFilePointer_beg
mov word ptr es [ sttFilePosition_low ][ di ], dx
mov word ptr es [ sttFilePosition_high ][ di ], cx

; Return
;-----
_moveFilePointer_36
RetCallersStackFrame ds, bx
mov word ptr [ _AX ][ bx ], dx
mov word ptr [ _DX ][ bx ], cx
Return

```

If you want a disassembly of genuine MS DOS but don't want to DIY, do it yourself, and for some time you will be a guy with a disassembly of DOS 1.1 or 2.1. Information Modes-Diction GX says ways to disassemble Int21 of these various of DOS. Imodes used the information given from its long ago disassembly project as part of its well known product, The 525 Network Scripta. We make recovery. Over 15,000 sold. For example, Figure 6-25 shows Imodes' rendition of the Crc and Set PSP functions from DJVMM, a disassembly dated April 1987. It is an inter-

existing reflection on the state of knowledge about DOS internals at the time that function 24 is labelled "get device driver list".

**Figure 6-25: Inodes Disassembly of DOS 2.1 Set and Get PSP**

```

J..... Set current PSP ..... Fn 50
L1006:
MOV CS:L0191,BX ;current PSP seg
RET_NEAR

J..... Set current PSP ..... Fn 51
L100C:
CALL L0C1A ;ds:si--> user's stack
PUSH CS:,0191
POP [SI+2] ;return in bx
RET_NEAR

```

Figure 6-26 shows the Inodes interpretation of the seek function from DOS 2.1, which can be compared against the MS-DOS 6.0 disassembly in Figure 6-20 and the BvDOS implementation in Figure 6-24.

**Figure 6-26: Inodes Disassembly of DOS 2.1 INT 21h AH=42h (lseek)**

```

J..... lseek (handle) ..... Fn 42
;bx handle
;cx,dx = hi_low dword offset
;al seek mode, 0 = from file start
;          1 = from current position
;          2 = from file end
;return cx=0, dx_ax = new position (from start)
;        = 0" -
;return, cx!=0, ax = 1 = invalid function (mode)
;        = 6 = invalid handle

L3B05:
CMP AL,3 ;is method in range 0..2 ?
JC L3B0D ;no; yes-->
MOV AL,1 ;err = invalid function

L3B0B:
JMP SHORT L3B03 ;do a error return

L3B0D:
PUSH SS
POP DS
CALL L3BF0 ;with bx=handle, get handle defn
PUSH ES
POP DS
JC L3B01 ;if handle bad--> ret, invalid handle
TEST BYTE PTR [DI+10h],80h ;is char device?
JZ L3BF2 ;yes; no-->
XOR AX,AX ;record = 0 always
XOR DX,DX
JMP SHORT L3C08 ;--> set random record fields

L3BF2:
DEC AL ;is method 0, from file start ?
JL L3C05 ;no; yes-->
DEC AL ;is method 1, from current position ?
JL L3C18 ;no; yes >

J..... method 2, from end of file
XCHG DX,AX ;ax = LSWord
XCHG DX,CX ;dx = RSWord
ADD AX,[DI+13h] ;add fcb's file size
ADC DX,[DI+15h]
JMP SHORT L3C08 ;--> set fields

```

```

;..... method 0, from start of file
L3C05.
    XCHG DX,AX          ;ax = 15Word
    XCHG DX,CX         ;dx  = 15Word
    
```

As with the PSP function, this disassembly of `back` in DOS 2.1 bears many similarities to the disassembly of `back` in DOS 6.0. On the other hand, the DOS 2.1 version does not do Windows and does not contain any network redirector code.

### Microsoft's DOS OEM Adaptation Kit (OAK)

But perhaps you are, deep and desperately, about getting the genuine article—commented source code—for Microsoft's MS-DOS 5.0 and higher. Microsoft does not publicize the product a great deal, but Microsoft will sell you an OEM Adaptation kit on signing a license agreement. Microsoft's OAK comes in an oddily formatted tape cartridge, but a version on normal PC diskettes is available from Annabooks, San Diego, CA.

The contents of the OAK are Microsoft confidential, so unfortunately we cannot reproduce any of it here, but we can give you some idea of its contents:

```

\BOSSDOAK
  \BIOS
    msbio1.asm
    msbio2.asm
    system11.asm
    system12.asm

  \BOOT
    msboot.asm

  \CMD
    \COMMAND
    \FORMAT
    \MODE

  \DEV
    \ANSI
    \HIREM

  \DOS
    fat.obj
    getset.obj
    handle.obj
    msdisp.obj

  \M
    cds.h
    dpb.h
    sysvar.h

  \INC
    arena.inc
    bpb.inc
    mult.inc
    pdb.inc
    sysvar.inc
    win386.inc
    wpatch.inc
    
```


As you can see from this very partial directory tree, Microsoft supplies some components of the OAK in ASM source code form, and others are supplied as OBJ files. The idea, of course, is that the OEM will change parts of IO.SYS but not MSDOS.SYS, so IO.SYS comes with source, but MSDOS.SYS









**Figure 7 1. Organization of the DOS Memory Arena**

The total memory structure is referred to in other documentation—and in Ray Duncan's better than official *Internal MS-DOS Programming*—as the memory arena, and the MCB that begins each block is called an arena header. Throughout this column, though, we refer to the arena header by the term MCB.

The entire MCB structure is built at system boot time, just after the parsing of CONFIG.SYS directives. This structure sits in memory below the DOS data segment because all RAM in that area was assigned early in the boot-up procedure and is not subject to reallocation.

### Memory Control Blocks

Each block of memory begins with an MCB, which is a single paragraph. That is, the MCB is 16 bytes long and begins at an address that is an exact multiple of 16. Memory blocks themselves are also always an exact number of paragraphs in length. This paragraph alignment makes it possible to refer to a memory block using a 16-bit segment address rather than a full 20-bit address.

The exception to the VNTDK DOS-11 header rule, which we provide to supply full documentation in one place, of the various undocumented structures our sample code deals with, it's on the companion diskette: in Figure 7 2 shows how each MCB is organized.

**Figure 7 2. DOS Memory Control Block Structure**

```
typedef struct {          /* Memory Control Block entry */
  BYTE type;             /* 'R' in chain; 'Z' not used */
  WORD owner;           /* PSP of the owner */
  WORD size;            /* in 16-byte paragraphs */
  BYTE unused[3];
  BYTE owner_name[8];   /* filename of owner, w/ DOS4 */
  } MCB, far * LPMCB;
```

When a program requests a block of memory with INT 21h function 48h, DOS must find the number of requested paragraphs, plus one more for the MCB. Assuming that a block of memory is available, DOS sets up its first paragraph as an MCB and hands the segment address of the second paragraph back to the program. Let's say you've made this call:

```

mov ah, 48h          , Allocate Memory Block
mov bx, 1           , one paragraph
int 21h
jc fail
, AX now holds initial segment of allocated block

```

Let's say AX now holds the value 1234h. This means that an MCB is located at 1234h. What does this MCB look like?

The first byte of every MCB (MCB type, except the last one in the chain) is 4Ch, the ASCII code for M. The next MCB's first two bytes instead 5Ah / Z. It may only be coincidental that these two letters are the initials of a well-known architect of the DOS memory manager, Mark Zubowski. In our example, the next could contain M or Z, enough M's to form more files.

Following this tag, we see a 16-bit value, a 16-bit word high format MCB pointer, that identifies the owner of the MCB. In fact, it will be 0000 (the segment block is available to use). Otherwise, it will contain the ID number of the process to which the block has been allocated. This word pointer. This information is used to locate free blocks, those where MCB pointers to, and to release allocated blocks when a process terminates. This ID number is the Program Segment Prefix, see below, of the user, in our example, this field would hold the PSP of which program was the current process when INT 21h function 48h was called.

Following the owner word is another word, MCB size, giving the size in paragraphs of the memory block controlled by this MCB. In our example, this field will be set to 1, indicating that the MCB at 1234h controls only the next paragraph at 1234h. In other words, the size value does not include the paragraph taken by the MCB itself, consequently, it is possible to have a valid MCB that shows a free block of size equal to 0. This happens when a full paragraph of a previously free block is allocated, instead of a whole paragraph. Because the MCB is some number of paragraphs, not the number of bytes, it can manage blocks up to 1048576 bytes in size, which is the maximum space available to a block. This is the basis for huge programs, such as DOS.

The last bytes of the word are unused applications of MS-DOS to date. In versions 2.0 and 3.0, the first two bytes of the MCB were unused, since DOS 4.0 and later versions the final eight bytes of a MCB (MCB low number) may contain the fragment of the covering program. For very small programs, the MCB then controls the memory used by the program's PSP; otherwise, the final eight bytes are ignored.

To find the next MCB in memory, consider it is not a linked list. To start with, the MCB's own segment address will have to give the segment address of the RAM it controls, then add to that the size to obtain the location of the next MCB. The reason is the segment address of the next MCB. For example, the next MCB is at 1238h. If the next at 1238:0000 is anything other than 'M' or 'Z', the MCB chain has been interrupted and continuation operations are not possible.

The first MCB always the one that controls DOS's own 512K segment, this MCB contains the memory address of the BIOS commands given as CONFIG.SYS. In some more recent versions of DOS that I have seen, this is 00000h or 0 to reason that I have ever seen about a sector.

The final MCB is determined by the /w or its first byte and, in non-matched 40k systems, generate a next MCB address of 0A0000h, which that address should not be used because the /w will include 2Ks for next MCB systems. Under DOS 5.0 or word card versions, when using a third party memory manager such as QUAM or MEMMAX, it is possible to enable Upper Memory Blocks. When this is done, the MCB chain may extend past segment 0A000h, since the point that ROM BIOS is found. We look at this later in a separate section.

In DOS 4.x, a 16-bit word, the DOS data segment memory block that is the RAM that DOS retains for its own use, cannot be allocated using the first of as is described later in this chapter. In spite of this, in consequence, each subsegment has its own variant of the standard MCB. However, the MCB's next MCB pointer field segment includes the entire memory segment, not only the subsegment which it points through the MCB chain.

These subsegments follow a format similar to, but not identical with, the MCB layout. The first byte is a 00H indicating 5256. In the word at byte 1 is not the owner. Instead, it is the actual segment address of the control block. The word at byte 3 is the size in paragraphs of the control block. Bytes 5 through 15 contain the filename, padded with blanks, of the file from which a driver was loaded.

Table 7-1 shows the codes used in these subsegment control blocks.

**Table 7-1 Codes Used in Subsegment Control Blocks**

Code	Directive Meaning Device Driver, if Present
E	device driver appendage, if present
I	IFS- installable file system driver, if present
F	FIFS- control block storage area (for FIFS 5)
A	FIFS- control block storage area, if present
C	FILESYS-MS work space area (BUFFERS- A option used)
B	BUFFERS- storage area
L	LASTDRIVE- drive info table storage area
S	STACKS- code and data area
T	Transient area (DOS 5 and later only)

In a control block, the device driver blocks appear only in the DOS data segment, a data meaning only a device driver appears to be documented. One other purpose appears to be simplification of the MEMMAN record in the DOS 4.0 through 5.0. Most of the information contained in them is duplicate information available from applicable DOS structures.

Some file allocations are done in the buffer used internally by DOS for parsing CONFIG.SYS. A space is reserved for a group and D represents DRIVE. The letters at one are not identical because they map to the CONFIG.SYS buffer block to represent STACKS since S apparently would be the MHI1-2-3-4-5. For more information on the CONFIG.SYS buffer, see Microsoft's *Microsoft Windows CONFIG.SYS Options at Boot*, *PC Magazine*, 29 November 1988, pp. 323-34. A more detailed explanation of the DOS files.

### The HMA and UMBS

Below we examine the details of using the MCBs in the next step, and look at two particular areas of memory management in DOS 5.0. These are the High Memory Area and Upper Memory Blocks.

When most computers manufactured RAM stops starting address A0000000. This address is the start of the 80386-16MMS segment, which is reserved for use with video RAM and is otherwise unavailable to users and the OEM BIOS. Also, memory space addressable by even the original 80286 processor was limited to up to address FFFF0000.

Normally, after an 80386 processor appeared, responsibility to a few more than a megabyte of address programming was needed to make use of parts of that 384k that were now taken up with video RAM and system services. Some before DOS itself had a way to deal with these areas, but post-80286 computers such as QEMM and 386MAX provided capabilities to "load programs right in the middle of the 384k that were not otherwise used."

Then Microsoft developed its AMS specification that dealt with using the High Memory Area. A group of users by University Park Subsegment. Microsoft needed provision for allocating memory in upper memory. In other words, for the 384k address intended area. With the release of DOS 5.0, allocation of UMBS became officially recognized and documented. Figure 7-7 is a memory map summarizing where these areas fit with respect to conventional RAM.

**Figure 7-3. Memory Map Showing UMB Region and HMA**

### Making Use of UMBs

To make use of Upper Memory Blocks, it's necessary to include a special command (DOS UMB) in CONFIG.SYS to enable them. You must also include EMM386.EXE (with its option to use RAM) or NUMA.SYS (or some other memory manager that's a UMB provider) to take its place. Some third-party memory managers don't require the DOS UMB line; check your manual for details.

With these prerequisites out of the way, all you do is use the internal command `UMB,MBL1` (or its alias `LF`, processed via an internal program command file). This command, used primarily when loading TSRs, causes the command interpreter to link the UMB memory chain to its normal conventional RAM-MB linkage. The command interpreter also sets the allocation strategy so that DOS searches the UMBs first when looking for space to load your programs into. If DOS cannot find enough space in Upper Memory, conventional RAM will be used, without error or warning message. When your program returns to DOS, the allocation strategy is returned to its original state and the UMB chain is unhooked from the conventional-MB linkage.

The DOS actions involved were documented in the *Commander's Reference to MS-DOS 5.0* (the method by which `UMB,MBL1` is implemented, however, has not been). The simple INTRSY script in Figure 7-4 provides a sequence of events.

**Figure 7-4. Organization of the DOS Memory Arena**

```

; LOADHIGH,SCR
intercept 27h
  on_entry
    output I21 at " cs: " ip ", AX=" AX ":
  function 0Ah
    on_entry
      sameLine "Buffered input to " ds ":" dx
  function 25h
    on_entry
      sameLine "Set INT" at " ==> " ds "1" dx
  function 26h

```

```

    on entry
        sameline "Create new PSP at " dx
function 30h
    on_entry
        sameline "Get version, flag = " al "
    on_exit
        sameline "Ver " al ". " ah ", 00x bh
function 35h
    on_exit
        sameline "INT" al " adr is " es ". " bx
function 37h
    on_exit
        sameline "Switcher is " dl
function 48h
    on_entry
        sameline "ALLOC " bx "h paras"
    on_exit
        if (cflag=1)
            sameline " FAIL (" es ")", only " bx "h available"
        if (cflag=0)
            sameline " returned seg " ax "h"
function 49h
    on_entry sameline "FREE seg " es "h"
    on_exit (1 (cflag=1) sameline " denied (" es "h)"
function 4ah
    on_entry
        sameline "REALLOC seg " es "h to " bx "h paras"
    on_exit
        if (cflag=1)
            sameline " FAIL (" es "h)", only " bx "h available"
function 4bh
    on_entry sameline "Loading " (dx dx+byte,ascii,64)
    on_exit output "----ret from child process----"
function 50h on_entry sameline "Set PSP, " BI
function 51h on_exit sameline "Get PSP " BI
function 58h
    subfunction 00h
        on_entry
            sameline "Get Allocation Strategy"
        on_exit
            sameline " returned " ax
    subfunction 01h
        on_entry
            sameline "Set Allocation Strategy to " bx
        on_exit
            if (cflag=1)
                sameline " FAILED (" es "h)"
    subfunction 02h
        on_entry
            sameline "Get UPB Link"
        on_exit
            sameline " returned " al
    subfunction 03h
        on_entry sameline "Link/Unlink UPB to PCBs: " bx
        on_exit
            if (cflag=1)
                sameline " FAILED (" es "h)"

```

When this script compiles and runs under IBM386 and COMMAND.COM, here's the essential part of the resulting report. The command traced was "ECHOADIRGHI DIDIT" (DIDIT.COM is a tiny .COM file that simply outputs one line and terminates).

```

121 at 9672 59E4, AX=5800. Get Allocation Strategy returned 0000
121 at 9672 59E8, AX=5802. Get UPB Link returned 00

```



```

121 at 9672 596f, AX 5800 Set Allocation Strategy returned 0000
121 at 9672 597a, AX 5801 Set Allocation Strategy to 0080
121 at 9672 5983, AX 5803 Link/Unlink UMB to MCBs 0001
121 at 9672 29ff, AX 4903 FREE seg 1400h
121 at 086d 010c, AX 4800 Loading D:\WB032\ISP\YIDIT.COM
121 at c0e8 0107, AX 0900
---ret from chkd process---
121 at ffff 880e, AX 5800 Set Allocation Strategy returned 0080
121 at ffff 881e, AX 5801 Set Allocation Strategy to 0000
121 at ffff 8829, AX 5803 Link/Unlink UMB to MCBs 0000

```

You can see that the BIOS calls to function 58h came from the transient portion of COMMAND.COM in segment 5972. With the UMB region linked and the strategy set to allocate space in the UMB region first on a first-in basis, the resident portion of COMMAND.COM then invokes the DOS sub-allocation sequence (0F4D) then invokes the DOS loader.

The final trace call to function 09h from segment 0105 is the single line report sent by DIBH.COM; the trace confirms that the program actually did load into a 640-KB CNP way (0E801d) in the Upper Memory above 40000h. DIBH then terminated, returning control to the HIMM at segment 0114. There, COMMAND.COM reset the strategy to normal first fit, and DOS removed the UMB region link, completing the LOADHIGH sequence of events.

### The High Memory Area

Like UMBs, the High Memory Area can be used only with an 80286 or later processor. Its operation depends on the CPU having more than 20 address bits. The original 8088-based systems had only 20 address bits available. If you tried to address a bit, a FFFF 0000, for example, the CPU would perform an overflow and generate an absolute address on the system bus of 10000000. That, however, is not the bus address; the CPU simply ignored the excess bits so that a address of this sort would resolve to 0000 0000.

With the 80286 appearing to have 24 address bits, this wrap-around stopped being automatic. So many programs depended on this wrap-around occurring, though, that system designers added hardware for the explicit purpose of making it continue to happen. Of course, the added circuitry had to be capable of being selected in order to conserve 386 operations on small 24-bit words.

When the wrap-around was disabled, it was then became possible to address 520,000 to 640 bytes of extra DRAM above the 1-megabyte mark without having to connect the processor out of real mode. By using a set of memory banks of 128K each (depending on the wrap-around) or taking place programs into four 1-megabyte 128-word 1-megabyte RAM, the same as conventional RAM.

That specification has been named the High Memory Area (HMA) and one of HIMEM.SYS's purposes is to control access to the HMA. But, the AMS specification and the DOS 5.0 *Programmer's Reference* consistently use the HMA, but is otherwise to come about identical points about the way it really works.

For the specification designers, it is that of a 20-line, which controls access to the HMA is turned off, and the program is loaded into the DOS is loaded high with the DOS HIGH command the A20 line is disabled. It can be enabled for the operation of a DOS service. In that case, the program can write directly to memory. Since the loader initially DOS repeatedly in the case of a program, it will write into RAM, the A20 line, as well as set by interrupt 13h. CONTOUR.SYS controls the DOS HIGH command. You can verify this using DIBH.COM; just issue the command of 0F40. To enable the HMA display on the screen, You can scroll through the whole thing by using repeated "d" commands.

There are A20 line-related access problems with programs that expect wrap-around to occur, of course. One such situation is the unpacking routine Microsoft's own linker originally included with any file that has a .EXE file. To reduce its size, according to Polytech Corp., a box of the strategy, DOSMAX.UMB, a program utility and a vector in the DOS 5.0 assembly, is a fix.

was only "Packed Exec corrupt" error message that began appearing everywhere shortly after the introduction of DOS 5.0 is directly due to the fact that the A20 line is enabled, and the original algorithm for寻址 depended on the segment wraparound effect to properly expand the compressed files.

Specifically, the packing technique used by LEXPACK is a form of run length encoding, the expansion method involves moving the packed copy of the program up in memory so that its end is at the first available location, then working backward toward the front of the program. Each time the algorithm could find a run-length count code, it would use a target address for the program's first byte available for count. All bytes from the start of the program up to the count code are then moved back so they start at that lower target address and the repeated run is inserted at the appropriate location, covering the run in its code. The original version of the unpacker calculated the offset of the new target address, moving it to the count at 32-bit from the adding that to the current address in 20-bit form. When segment wraparound is in effect, the final calculation produces a 20-bit result of 000000h, which is not what the expansion was successful. Otherwise, when the program happens to be located in the top 64K of the 1MB address space, the final result wraps into the High Memory area, which causes the corrupt message to be displayed. Once the problem was identified of course. Microsoft did a little bit of compression work, and versions of the linker produced since mid-1988 or so don't treat it as a "More popular packages, though, were built with the older linkers, so the error may be serious for some users." In fact, it popped up even to the surface during final testing of DOS 6.0, because some of the new system programs were linked with an old copy of the linker!

For a complete DOS API for x86, see the HVM, see Chapter 4, and the appendix entries for INT 21h AX:4A01h and 4A02h.

### How To Find the Start of the MCB Chain

The key to finding the MCB is in the undocumented DOS List of Lists, whose address is retrieved with INT 21h function 22h. Although the List of Lists returned by this function differs from the one used in DOS 4.0 to locate the MCB pointer's location is one of the very few items that is the same in DOS versions to date. It is always associated two bytes in front of the pointer returned in ES:BX, that is at ES:[BX-2].

The value associated is not actually not a pointer to the first MCB but its segment number, that of the DOS 16-bit segment memory block mentioned earlier. To use it as a pointer, you must provide an offset of 0000h.

The following assembly code fragment shows how to force the MCB pointer into ES:SI so that it may then be used to retrieve the key byte, the owner word, and the size word; this code only sets up ES:SI and does not retrieve the data.

```

mov     ah, 52h           ; Get List of lists
int     21h
mov     ax, es:[bx-2]    ; First MCB Segment Address
mov     es, ax
xor     si, si           ; force offset to be zero

;-----
xor     bx, bx
mov     es, bx           ; set ES:BX to 0:0
mov     ah, 52h
int     21h
mov     cx, es
or     cx, bx            ; is ES:BX still 0:0?
jc     fail             ; then Function 52h not supported
mov     ax, es:[bx-2]    ; First MCB Segment Address
mov     es, ax
xor     si, si           ; force offset to be zero
fail

```

The next code sequence then retrieves the key byte (the owner word) and the size word, respectively; the code assumes that PSP is unchanged from the preceding example.

```
mov  al, es:[si]      ; gets key byte, 'M' or 'Z'
mov  bx, es:[si+13]  ; gets owner word or 0000
mov  cx, es:[si+33]  ; gets size in paragraphs
```

For most applications, a code fragment in Figure 7-5 may be more useful. It can be used in Microsoft C 5.0 and higher, QuickC 2.0 and higher, Borland TurboC 2.0 and higher, and Bor and C++ 2.0 and higher.

**Figure 7-5 Get First MCB() Routine**

```
#include <dos.h>          /* use standard header file */
#include "undocdos.h"     /* use our standard header file */
                          /* to define MCB structure and */
                          /* MK_FP macro (make far ptr). */

_LPMCB Get_First_MCB( void ) /* locate first MCB, return ptr */
{ union REGS reg;        /* REGS, SREGS are defined by */
  struct SREGS seg;      /* the DOS H header file */
  WORD *tmp;

  segreg( &seg );        /* set up seg regs */
  reg.h.eh = 0x52;       /* get List of Lists in ES:BX */
  intdosx( &reg, &seg, &seg );
  tmp = (WORD far *) MK_FP( seg.es, reg.bx - 2 );
  return (LPMCB) MK_FP( tmp, 0 );
}
```

Get\_First\_MCB() is functionally identical to the first assembly language fragment. This function uses the MK\_FP macro, which became MK\_FP in Visual C++ 1, to create the far and pointer value, rather than stuffing the appropriate quantities into DS and SI. As explained in Chapter 7, you can also use inline assembly to register pseudo-variables; your compiler supports this option.

## How To Trace the MCB Chain

Let's look at how to locate, print, and trace through the MS-DOS MCBs and determine how RAM is allocated. You'll need a debugger, such as a program on your machine. A good one is this popular one: *ImagE MEM*—a standard set of DOS 4.0 and later TMAP. Chris Dunford, MAPMEM—Commercial Software, and DMMEM—Borland Turbo Debugger 2.0. However, our version of DMMEM is a computer-independent low-level utility we write.

Because some MCBs store PSPs, this program can be used to trace through all PSPs, showing which programs are resident in memory. When you refer to MCBs controlling PM, we only mean that the block of memory contains PM. MCB happens to be a program to be traced if not a program to be traced; a program may have been loaded into memory. All DOS processes begin with a 256-byte "free program" PSP. The MCBs control the PSP only in the sense that the MCB's byte are a necessary condition for use by the PSP and by the process itself. For example, a PSP at 0419 is controlled by MCB at 0408. If you know the owner field of the MCB at 0408 would be 0419.

Our DMMEM program loops through segments in memory of each MCB, the Program Segment Prefix of a program, and lists the MCB in hex paragraphs and decimal bytes. For MCBs that hold actual PSPs, DMMEM also lists the segment for the corresponding environment. The ASCII hexadecimal byte owner word in DOS 3.0 and higher is kept in program environment, and an interrupt vector table pointer at the end of memory. The program also shows subsegments within the DOS data segment, if you so desire, and knows about UMVs.

One feature of our MCB walker is that it can assemble the presence of one or more MCBs and, in fact, programs such as SPMMV and QEMM allow memory resident programs to be loaded. Each

by creating a secondary MCB chain in high DOS memory. With the advent of UMBs in DOS 5.0, this is the rule rather than the exception. Our UDMEM program shows secondary MCB chains. Figure 7-6 shows what UDMEM's output looks like.

Figure 7-6. Sample of UDMEM Output

```
D:\UDOS2\CHAP7> udmem
Seg  Owner  Size
0255  0008  092A ( 37536)  DOS Data Segment
-----
Seg  Size  Type
-----
0254  0041  Device Driver (386MAX)
0296  0015  Device Driver (386LOAD)
02AC  0741  Device Driver (5STORDRV) [26 F5 FA FE ]
09EE  0015  Device Driver (386LOAD)
0A04  0050  System File Tables
0A62  0005  FCBs
0A68  0020  Buffers
0A89  0037  CDS Table
DA47  00BC  Status [02 0A 0B 0C 0d 0E 7D 72 73 74 76 ]
087E  0008  0004 ( 64)  BIOS Code area
0881  0084  0010 ( 256)  Env at 0374  chap7 new /pd:\ud0s2\chap7\udkeys /g
0894  0084  0005 ( 80)
089A  089B  0147 ( 5232)  Env at 0895  ff 2048 /i:20
DCE2  0CF1  0000 ( 208)
DCFD  0CF1  4708 (154800)  Env at DCE5  D:\UTILS\PE2 EXE chap7.new /pd:\ud0s2\chap7\udkeys
/g [30 65 66 ff ff ]
59CC  5908  0008 ( 208)
59DA  5908  0010 ( 256)  Env at 0A56
59EB  59FB  0008 ( 176)
59F7  59FB  12A8 ( 75392)  Env at 59EC  D:\UDOS2\CHAP7\UDMEM.EXE [00 E9 F9 ]
6C6D  0000  339E (211424)  free [EE EF FA ]
9FFF  End of conventional RAM

UMB Chain
CB00  FFFF  058E ( 22752)  386LOADED Driver [13 15 2B ]
C8BF  FFFA  0004 ( 64)  386MAX UMB control block
C894  FFFE  0205 ( 8272)  386MAX UMB
C19A  FFFA  0004 ( 64)  386MAX UMB control block
CF9F  FFFE  0022 ( 544)  386MAX UMB
CFE2  FFFA  0004 ( 64)  386MAX UMB control block
CFE7  FFFE  0022 ( 544)  386MAX UMB
CFEA  FFFA  0004 ( 64)  386MAX UMB control block
CFE7  FFFE  0042 ( 1056)  386MAX UMB
D032  FFFA  0004 ( 64)  386MAX UMB control block
D037  FFFE  0210 ( 8448)  386MAX UMB
D268  FFFA  0004 ( 64)  386MAX UMB control block
D26B  FFFE  00C0 ( 3072)  ( 1404 AD05 [0R [2E ]
D30E  FFFA  0004 ( 64)  386MAX UMB control block
D313  FFFE  0020 ( 512)  386MAX UMB
D33A  D33C  0006 ( 96)
D37B  D37C  04C0 (19664)  Env at D80A  C:\UV\UV.COM [10 ]
D8D9  D33C  0002 ( 32)
D8DC  D814  0006 ( 96)  [20 21 27 29 ]
D813  D814  01AD (6864)  Env at D9C2  D:\UTILS\CTRLALT.COM [0B 09 ]
D9C1  D814  0002 ( 32)
D9C4  FFFA  0004 ( 64)  386MAX UMB control block
D9C9  FFFE  0066 ( 1632)  386MAX UMB [22 23 24 2F ]
DA30  FFFA  0004 ( 64)  386MAX UMB control block
DA35  FFFE  0020 ( 512)  386MAX UMB
DA56  0000  05A8 ( 23168)  free
DFF7  FFFD  1200 ( 73728)  386MAX locked-out area [EB ]
F2D0  FFFF  0059 ( 1424)  386LOADED Driver [7B ]
F25A  0000  05A4 ( 23104)  free
F777  FFFD  0400 ( 16384)  386MAX locked-out area
```

```

FC00      FFFF      00AB ( 2768) 386LOADED Driver [40 67 ]
FC0A      FCB6      0006 (   96)
FC0E      FCB6      0042 ( 1056) Env at FCF9   C:\404\KSTACK.COM [16 ]
FC12      FCB6      0002 (   32)
FC16      0000      0004 (   64) free
FDD0      End of UMB Chain

```

To create this sample of output, I shelled out to DOS from within my text editor program (D:\THIS IS THE FILE) at address 0011. The output shows both the normal RAM area below 0A000 0000 and the Upper Memory Blocks above that point. It also shows how various interrupts are routed to different drivers and programs, both above and below the UMB line. These references to such interrupts as FC 1E and FC 1F result from the fact that during the boot up process, initialization code uses a top of the interrupt vector area as its stack and never cleans the memory or other ward. These references are not valid.

It is useful to write UDMEM in two stages. First, just print out raw information about DOS memory control blocks. Then, after that simple program is working, write an improved version that displays the UMIB hierarchy of the owners of the UMIBs, which goes as a display of all programs resident in memory, including, of course, the UDMEM program itself. Figure 7-7 shows our first version of the UDMEM utility.

Figure 7-7. First Version of UDMEM

```

/*
UDMEM1 ( -- walks DOS MCB chain(s); simple version
Andrew Schuman and Jim Kyle, July 1990
revised by Jim Kyle, August 1992, March 1993
*/
#include <stdlib.h> /* needed by RSC only */
#include <stdio.h>
#include <dos.h>
#include "undocdos.h"
void fail(char *s) { puts(s); exit(1); }

LPCB get_mcb(void)
{
    ASM mov ah, 52h
    ASM int 21h
    ASM mov dx, es:[bx-2]
    ASM xor ax, ax
    /* in both Microsoft C and Turbo C, far* returned in BX AX */
}

void display(LPCB mcb)
{
    char buf[80],
    sprintf(buf, "%04X %04X %04X (%04u)",
    FP_SEG(mcb), mcb->owner, mcb->size, (long) mcb->size << 4),
    if (! mcb->owner)
        strcat(buf, " free"),
    puts(buf),
}

void walk(LPCB mcb)
{
    printf("Seg Owner Size\n");
    for (;;)
        switch (mcb->type)
        {
            case 'H' : /* Mark : belongs to MCB chain */
                display(mcb);
                mcb = (LPCB)PK_FP(FP_SEG(mcb) + mcb->size + 1, 0);
                break,

```

```

case 'Z': /* Zbikowski : end of MCB chain */
    display(mcb);
    printf("Z044\n", FP_SEG(mcb) + mcb->size + 1);
    return;
default:
    fail("error in MCB chain");
}
}

```

```

main(int argc, char *argv[])
{

```

```

    walk(get_mcb());
    //could walk UMB chain here too
}

```

IBM PC DOS 3.31 simply displays the raw MCB chain and makes no attempt to connect to any upper memory blocks that may be present. The function `get_mcb()` written with inline assembler (ctor is a fair guess) searches the MCB list through what is a long and somewhat tedious DOS function 52h here, we don't bother to use IBM's version of `subsys` because the segment of the first MCB is always located at offset 2 in the list of lists. It is even supported in the DOS compatibility boxes of OS/2 and Windows NT, see Chapter 4. The start of the MCB chain is passed to the function `walk()`, which goes into a `while` loop displaying an MCB and moving to the next MCB until the end of the chain is reached or a failure occurs. The MCB is displayed using the function `display()`. The output of this program looks like this:

Seg	Owner	Size	
09F3	0006	03E1 ( 15888)	
0005	0006	0003 ( 3376)	
0E49	0000	0003 ( 48)	free
0E40	0006	0040 ( 1024)	
0EE8	0006	0004 ( 64)	
0EF3	0F02	0000 ( 208)	
0F01	0F02	1204 ( 75792)	
2106	0000	7EF9 (520080)	free
0000			

### MCB Consistency Checks

After this side's inclusion built. After chopping out `main()` it can be linked into other programs in the set to check the DOS memory allocation. This is particularly useful when you are trying to debug a program that trashes the MCB chain by modifying the `walk()` function into `mcb_chk()` as shown in Figure 8. You can check the MCB chain for consistency before DOS does. The chain is correct if it does not equal to anything other than `MEM0 /`.

### Figure 7 8. Routine to Check MCB Chain

```

BOOL mcb_chk(LP_MCB mcb) /* see UNDISCOVERED.DOS for typedefs */
{
    for (;;)
        if (mcb->type == 'M')
            mcb = LP_MCB_PTR(FP_SEG(mcb) + mcb->size + 1, 0);
        else
            return (mcb->type == 'Z');
}

```

With `mcb_chk()` a program can periodically check the MCB chain with a call such as the following:

```

if (! mcb_chk(get_mcb()))
{
    /* maybe do stack backtrace here, or dump registers */
    puts("Error in MCB chain - prepare for halt...");
    getch();
}

```

Of course, `malloc` does return `NULL`. But the next time any memory allocation is performed, the system will halt with a message such as

```
Memory allocation error
Cannot load COMMAND, system halted
```

DOS internally performs the same consistency check as `malloc`—except that, if it does find anything other than `MEM_0` or `Z`, DOS has no choice but to halt the system. There seems to be no way that the MCB chain could be reliably repaired. In multitasking 80386 control programs such as DEMView or Windows 3.x, you might trash the MCB chain in a DOS box, virtual machine, or fat-less catastrophe. You just throw the entire machine away and get a new one.

Our original MCB loader has one other use. We can use it to reveal a bug in DOS itself. If the entry for `INT_21:FUNCTION_4Ah` (Resize Memory Block) in the appendix notes that “if there is sufficient memory to expand the block as much as requested, the block will be made as large as possible.” Don’t believe it. Just substitute the version of `main` shown in Figure 7-9 for the one provided in Figure 7-7.

### Figure 7-9. How To Demonstrate a DOS Bug

```
main(void)
{
    unsigned segm,
    ASM mov ah, 48h          /* Allocate Memory Block */
    ASM mov bx, 04h        /* get 100 paragraphs */
    ASM int 21h
    ASM jc done
    /* ax now holds initial segment of allocated block */
    ASM mov segm, ax
    printf("before "); display(MEM_FP(seg - 1, 0));

    ASM mov ax, segm
    ASM mov es, ax         /* now resize the block */
    ASM mov ah, 4Ah       /* Resize Memory Block */
    ASM mov bx, 0FFFFh    /* impossible in real mode */
    ASM int 21h
    ASM jnc done          /* something badly wrong if didn't fail! */
    printf("after: "); display(MEM_FP(seg - 1, 0));
done
    return 0;
}
```

The resulting display shows that all remaining memory has in fact been given to the block—even though the call failed:

```
before: 104C 0BEA 0064 ( 1600)
after:  104C 0BEA 9A03 16356481
```

The entries `0BEA` and `0064` are assigned to MCB `104C` in the second line; shows that even though `FUNCTION_4Ah` failed with the carry flag set indicating an error, the block was still made as large as possible. It's particularly interesting because this test was run on a system with Quarterdeck QEMM. This system includes, e.g., DOS, not an undocumented feature on which you would depend! As it stands, allocations that fail but that nonetheless snag memory can cause system-as-program behavior.

The `system` case is that DOS assigns all free space to the program before discovering that there's not enough, but then fails to restore the original program allocation and put the free space back once the error has been detected.

This `main` uses `display` in the `display` function can be useful all by itself. Just pass the function an MCB and it displays some information. Given the segment address of a block of memory

to be able to determine that the MCB is located at the preceding paragraph. If a PSP is, for instance, 1234h, the MCB is 1234h. This is why `segment 1` rather than `segment 0` is used above in the call to `display()`.

### A More Detailed UDMEM Program

For `display()` to be able to specify `display`, we need to change the `display` function, hook into any UMB chain that may be present, and add supporting functions.

Since the relationship between MCB, PSP, and environment can get a little confusing, so we are going to step through the details that are defined in `UDMEM.H`:

```
#define MCB_ENV_SEG(segment) ((long) 1)
#define IS_PSP_mcb ((FP_SEG(mcb) + 1 == (mcb->owner) && \
    (WORD FAR *) MCB_FP_FP_SEG(mcb) + 1, 0) == 0x20CD))
#define ENV_M_PSP(psp, seg) ((WORD FAR *) MCB_FP_FP_SEG(seg, 0x2c))
```

`ENV_M_PSP(MCB_ENV_SEG(segment))` simply converts the PSP segment value to the corresponding MCB segment value. The last, `ENV_FP_PSP()`, grabs the segment address of the environment block in the MCB. The `IS_PSP()` macro performs a test on a segment value to determine whether it is a PSP segment. The first test simply verifies that this block owns itself. However, it is not a test of protected mode programs such as Windows 3.x and some recent DOS programs because they do not use such SPS as a test to indicate they are in a segment. Thus if a segment passes the first test we go on to verify that the first two bytes of the segment are 0120, the machine code for INT 20h. All of this is done so we have two logical posters that we have can into so far fail the test.

### Figure 7-10. New Display Routine

```
void display1(MCB mcb)
{ static FP_vect_2e = { FP } 0,
  WORD env_seg;
  p = " 04x 104x 204x (161u) ",
    FP_SEG(mcb), mcb->owner,
    mcb->size, (long) mcb->size << 6 }

  if (IS_PSP(mcb))
  { FP = env_seg; /* MSC wants lvalue */
    if (env_seg == FP_SEG(e)) != 0)
      printf("Env at 104x ", env_seg),
        printf("No Env Segment ");
  }

  if (vect_2e)
    vect_2e = GETVECT(0x2e); /* do just once */

  /* INT 21h be using master COMMAND.COM for other shell */
  if (getenv("vect_2e", FP_SEG(mcb), mcb->size))
    printf("2e ", getenv("COMSPEC"));
  switch (mcb->owner) /* decode special stuff */
  {
  case 0:
    printf("free ");
    break;
  case 6:
    printf("DR-DOS XMS UMB ");
    break;
  case 7:
    printf("DR-DOS hole ");
    break;
  case 8:
    printf("DOS ");
    display_subsegs(mcb, mcb->size);
  }
```



```

    return, /* display_subsegs cleans up */
    case 0xffff:
        printf( "386MAX UMB control block " ),
        break,
    case 0xffffd:
        printf( "386MAX locked-out area " );
        break,
    case 0xffffe:
        printf( "386MAX UMB " ),
        break,
    case 0xfffff:
        printf( "386LOADED Driver " ),
        break,
    }
    display_progname( mcb ), /* moved from original location
    if( IS_PSP( mcb ) )
        display_cadline( mcb ),
    display_vectors( mcb );
    printf( "\n" );
}

```

The new `display_umb` (Figure 7-10) calls `env` (Figure 7-11) to find out if the MCB contains the PSP of its owner and therefore has an associated environment block.

### Figure 7-11. Environment-Locating Routine

```

char far * env( LPMCB mcb )
{
    char far * e,
    WORD env_mcb,
    WORD env_owner,

    /* if the MCB owner is one more than the MCB segment then
    *   psp = MCB owner
    *   env_seg = make_far_pointer( psp, 2ih )
    *   e = make_far_pointer( env_seg, 0 )
    * else
    *   return NULL
    */
    if( IS_PSP( mcb ) )
        e = MAKE_FAR_PTR( ENV_FROM_PSP( mcb->owner ), 0 ),
    else
        return ( char far *) 0,

    /* Does this environment really belong to this PSP? An
    * environment is just another memory block, so its MCB is
    * located in the preceding paragraph. Make sure the env
    * MCB's owner is equal to the PSP whose environment this
    * supposedly is! Thanks to Rob Adams of Phar Lap Software
    * for pointing out the need for this check, this is a
    * good example of the sort of consistency check one must
    * do when working with undocumented DOS.
    *
    * Note that with DOS5, this test had to be changed to just
    * reject free MCBs, because 386MAX added special codes in
    * the owner field that caused our original test to fail!
    */
    env_mcb = MCB_FROM_SEG( FP_SEG( e ) );
    env_owner = (( MCB far *) MAKE_FAR_PTR( env_mcb, 0 ))->owner,
    /*return ( env_owner == mcb->owner ) ? e : ( char far *) 0; changed! */
    return ( env_owner ) ? e : ( char far *) 0;
}

```

The `env` function uses the `IS_PSP` macro to verify that this MCB is a PSP and that the `env` routine pointer is not NULL. In its original version, `env` further made sure we didn't pick up a

means of environment for a program that has freed its environment. Usually, such programs don't bother to zero out the environment segment number located at offset 2Ch in the PSP. The next-to-last byte of the program's own command-out was the original test. After Quantum's 386MAX memory manager modification, the owner field of the MCB for other purposes, it was necessary to change this function.

The display routine interprets the owner field of the MCB to decode all the special values that Table 7-11 appears in. If this shows the block to be the DOS data segment, display calls display\_subseg() (Figure 7-12) to list the subsegments.

**Figure 7-12. Decoding DOS Subsegments**

```
static BOOL sadone = FALSE;

void display_subseg( LPRCB mcb, WORD subsegz )
{ char tmp[ 10 ],
  char src = src;
  int i;
  if( sadone ) /* if not first DOS area */
  { printf( "Code area\n" );
    return;
  }
  printf( "Data Segment\n" );
  if( _osmajor < 4 ) /* subsegments only in V4+ */
  return;

  printf( " Seg Size Type Init\n", "-----\n" );
  mcb = MK_FP( FP_SEG( mcb ) + 1, 0 );
  while( ( int i = subsegz > 0 ) /* process each subsegment */
  { printf( " %04x %04x ", FP_SEG( mcb ), mcb->size );
    src = mcb->owner_name[ 0 ]; /* copy filename for B, I */
    if( i < 8 || src > 'Z' ) i = 0;
    tmp[ 1 ] = src;
    tmp[ 2 ] = 0;
    switch( mcb->type ) /* translate type codes */
    {
      case 'D':
        printf( " Device Driver (%s)", tmp );
        break;
      case 'E':
        printf( " Device Driver appendage ",
        break;
      case 'I':
        printf( " IPS Driver (%s)", tmp );
        break;
      case 'P':
        printf( " System file Tables ",
        break;
      case 'G':
        printf( " PCBs ",
        break;
      case 'C':
        printf( " Buffer Workspace ",
        break;
      case 'B':
        printf( " Buffers ",
        break;
      case 'L':
        printf( " CDS Table ",
        break;
      case 'S':
        printf( " Stacks ",
        break;
      case 'T':
        printf( " Transient Code ",

```



**Figure 7 14 Showing Command Arguments**

```

void display_cmdline( LPRCB mcb )
{ /* psp := MCB owner
 *  cmdline_len := psp[50h]
 *  cmdline := psp[81h]
 *  print cmdline(display width = cmdline_len)
 *
 * int len = *((BYTE far *) MK_FP( mcb->owner, 0x50 ));
 * char far * cmdline = MK_FP( mcb->owner, 0x81 );
 * printf( "%2.0Fs", len, cmdline );
 */
}

```

Note the `fprintf` cmdline uses the `%s` print mask to display a far string, using the macro `MK_FP( seg, off)`. The variable `len` whose value may be zero. Sometimes garbage is printed by MEMM, or by any similar program, because the disk transfer area located inside the PSP overlaps the beginning of the system file.

The `belongs` macro belongs (Figure 7 15) determines if an interrupt vector points into the block of code `word` is in MCB. The `WORD`, `FP`, and `WORD` types are defined in `UNDOS.DOS.H`.

**Figure 7 15 Testing Interrupt Vectors**

```

BOOL belongs( FP vec, WORD start, WORD size )
{ WORD seg = FP_SEG( vec ) + ( FP_OFF( vec ) >> 4 ); /* normalize */
  return ( seg >= start ) && ( seg <= ( start + size ) );
}

```

Finally `display` calls `display_vectors` (Figure 7 16) to show any interrupts hooked by the program whose PSP is contained in this MCB. The function binds these hooked interrupts simply by setting CS:IP for the interrupt handler falls within its MCB.

**Figure 7 16. Showing Interrupt Vectors**

```

#define _IRRLOC
#define GETVECT(x)    _getvect(x)
#define _getvect(x)  _dos_getvect(x)
#define _dos_getvect(x)
#endif

void display_vectors( LPRCB mcb )
{ static FP * vec = ( FP *) 0;
  WORD vec_seg;
  int i;
  int did_one = 0;
  if ( ! vec ) /* one-time initialization */
  { if ( ! vec = calloc( 256, sizeof( void far *) ) )
    { fprintf( stderr, "insufficient memory" );
      return;
    }
    for( i = 0; i < 256; i ++ )
      vec[ i ] = GETVECT( i );
  }
  for( i = 0; i < 256; i ++ )
  { if( vec[ i ] && belongs( vec[ i ], FP_SEG( mcb ), mcb->size ) )
    { if( ! did_one )
      { did_one ++;
        printf( "\n", i );
      }
      printf( "-02x ", i );
      vec[ i ] = 0;
    }
  }
  if( did_one )
    printf( "\n" );
}

```

In DOS 4.0 and higher, some memory-resident software can be loaded using the `INSTALL` statement in `CONFIG.SYS`. Such programs can show up in the UDMEM space as MCBs that aren't associated with any program at all, which may have hooked interrupt vectors. Note that UDMEM can display vectors for all MCBs, even when there seems to be no associated program. For example, in DOS 4.0 and higher, if `CONFIG.SYS` contains the statement `INSTALL C:\EDOC\EDOC.M` to load Chris Dunsford's `EDOC` command-line editor, then UDMEM displays something like the following:

```
0EB1 0EB2 065F ( 26096) (1B 21 61 )
```

Another bunch of strange display vectors for all MCBs that occasionally we had orphaned interrupt vectors that point into free memory:

```
2A2A 0000 7 (4B2640) free (30 54 F5 FB )
```

`INT 30h` is a trap/jump instruction, not an interrupt vector, but `INTs 14h`, `15h`, and `18h` are real interrupt vectors. Let's hope no program invokes them while they're pointing into free memory! Since these vectors are used as the stack of user boot-up, it is essential that any program that tries to invoke them be certain they have first been set to valid interrupt service routines.

The string length function shown in Figure 7-17 lets us easily get the length of fat strings, even from a small model program.

**Figure 7-17. Determining String Length**

```
WORD fstrlen( char far * s )
{
  #if defined( _MSC_VER ) && ( _MSC_VER >= 600 )
    return _fstrlen( s );
  #else
    WORD len = 0;
    while( * s++ )
      len++;
    return len;
  #endif
}
```

In addition to the changes made to the `display` function, it's also necessary to modify `walk`, as shown in Figure 7-18, so that it automatically includes any Upper Memory blocks that may be active.

**Figure 7-18. Revised Walk Routine**

```
void walk( LPMCB mcb ) /* walks chain displaying data */
{ printf( "Seg Owner $ize\n" );
  for( ; ; )
    switch( mcb->type )
    {
      case 'M' : /* MCB belongs to MCBchain */
        display( mcb );
        mcb = MC_FP( FP_SEG( mcb ) + mcb->size + 1, 0 );
        break;

      case 'Z' : /* Zbikowski : end of MCB chain */
        display( mcb );
        printf( "X04E\n" );
        FP_SEG( mcb ) + mcb->size + 1 );
        if( FP_SEG( mcb ) < 0xA000 )
          ( printf( "End of conventional RAM\n" ),
            mcb = firsthi(); /* try to link to UMB area */
            if( mcb == ( FP ) 0L ) /* No UMB MCB's found */
              return;
          );
    }
}
```

```

        printf("UMB Chain\n");
    }
    else
    {
        printf("End of UMB Chain\n");
        return;
    }
    break;
}

default:
    fprintf(stderr, "error in PCB chain");
}
}

```

If UMBs are not set first UDMEM is running from conventional RAM, there are actually two separate PCB chains: the first starts at the end of conventional RAM, while the other is entirely in the UMBs. Once the UDMEM call on DOS is used to set UDMEM, the two chains merge into one.

To accommodate both with the case of two chains, the actions performed when the program starts are a little special and changed. If the current PCB is below the top of conventional RAM (0x0000), the program needs to find the function to bridge to the UMB region. Otherwise, the second chain is to be used, and the address has been encountered.

The first of two functions bridges the link to any UMBs that might be present as a separate end of shared memory. This is done by itself and is incorporated into the UDMEM program. This function returns a pointer to the UMBs if one is located or a far pointer to the first PCB if the address is not found. This was written by Kim Kolonen, president of TurboPower Software, and is used in a number of his TSR.COM utilities. The TSR.COM utilities include MARK and RELEASE, which, with TSRV, can be included; they are available on CompaqServe and on many BBS systems.

**Figure 7.19. Finding the UMB Chain**

```

/*
 * FIRSTMI - utility to locate first UMB address
 *
 * adapted from:
 *   XMS PAS - unit of XMS functions
 *   Copyright (c) 1991 Kim Kolonen, TurboPower Software.
 *   MEMU_PAS - utility unit for TSR Utilities.
 *   Copyright (c) 1991 Kim Kolonen, TurboPower Software
 *   by Jim Kyle, with permission of Kim Kolonen. Thanks, Kim!
 *
 * These are part of Kim's TSR.COM utility that includes
 * MARK and RELEASE among others.
 */

#include <dos.h>
#include "undocdos.h"

static BYTE xmsinstalled( void ) /* true if XMS ok here */
{
    ASM mov ah,0x30;
    ASM int 0x21,
    ASM cmp al,3,
    ASM jae Check2F,
    ASM mov al,0,
    ASM jmp Done,
Check2F
    ASM mov ax,0x6370,
    ASM int 0x2F;
Done
    ASM mov ah,0,
}

```

```

static FP XmsControl(Addr_t void) /* gets far ptr to code */
{ ASM mov ax,0x4310;
  ASM int 0x2F;
  ASM mov ax,bx,
  ASM mov dx,es,
}

static FP XmsControl;

static BYTE AnyUMB( void ) /* verify that a UMB exists */
{ ASM mov ah,0x10;
  ASM mov dx,0xFFFF; /* force an error return in BL */
  ASM call dword ptr [XmsControl];
  ASM xor ah,ah;
  ASM mov al,bl; /* return error result */
}

LPNCB *rsth( void )
{ WORD Segment;
  WORD Size,
  WORD Mseg,
  BOOL Done,
  BOOL Invalid,
  LPNCB M,
  LPNCB N;
  LPNCB Retval = (LPNCB)0L;

  if( !XmsInstalled() ) /* UMBs are not possible */
    return Retval;
  XmsControl XmsControlAddr(), /* get adr to test UMBs */
  if( ! AnyUMB() ) /* UMBs are not possible */
    return Retval;

  ASM int 0x12; /* find top of conventional RAM */
  ASM mov cx,6;
  ASM shl ax,cx,
  ASM mov Mseg,ax; /* start the search there */
  Done = FALSE;

  while( M = (LPNCB)NK_FPC( Mseg, 0);
    if( M->type == 'N' ) /* may be an NCB, if any, must be 2 */
      { N = M,
        Invalid = FALSE,
        do /* determine whether valid NCBs */
          switch( N->type )
          {
            case 'R' : /* try for next NCB via linkage */
              << (WORD)(FP_SEG(N)+(N->size)+1) <= (WORD)0xFFFF )
                N = (LPNCB)NK_FPC( FP_SEG(N)+(N->size)+1, 0);
              else /* can't be valid, keep looking */
                Invalid = TRUE,
                break;
            case '2' : /* found end of chain starting at M */
              Retval = M,
              Done = TRUE,
              break;
            default : /* chain failed test, keep looking */
              Invalid = TRUE,
              } while( !Done && !Invalid );
          }
      }
  if( !Done )
    { if( Mseg < 0xFFFF )
      Mseg++,
    }
}

```

```

        else
            Done = TRUE,
    }
} wht.of 'Done }
return Retval,
/* either NULL or first NCB pointer */
}

```

The `LIST` function works on a brute force basis. It first verifies that UMBs might be present by testing for a UMB pointer, then for function segments (normal pointers). Moving to the final paragraph of memory in RAM, it then iterates and tests each subsequent paragraph aligned byte for the "M" type signature. Each time a signature is found, it goes into an inner loop that attempts to verify that this is really a UMB header. If the chain reaches a "Z" type byte without finding the original M byte's address, it goes on to test the next high MCB. If the inner loop fails, the outer loop resumes. If the inner loop finishes without finding an MCB chain, the function returns `NULL` as its value to indicate that no UMBs are present.

It would be possible to avoid the brute force approach and simply use the same chain linking technique for `LOADHIGH` itself. Unfortunately, UDMEM would not work with versions of DOS prior to 5.0 even if UMBs were not used, because the linking function did not exist before DOS 5.0. The `UDMEM` file allows `UDMEM` to work with older versions and also helps make the following interesting point:

Not only UDMEM, including AMS manages to follow the same rules. Both QEMM and EMM386 use a recursive algorithm that collects free UMB space into regions and collects the regions into a separate set of MCBs. This makes UDMEM search a little easier if loaded high than it does when run over low RAM, like those managers. The `UDMEM` function creates the region chain and, as a result, can find regions that are not found by individual UMBs.

When `LOADHIGH` sees `MS-DOS` bypass the chain of regions and links the individual UMBs via a direct connection to the original MCBs. UDMEM then shows the individual UMBs and does not display regions or other blocks.

Under DOS 5.0, this is somewhat an interesting oddity, with DOS 6.0, it's more significant since the `LOADHIGH` command of DOS 6.0 allows you to specify which region to use for the program being loaded.

We can take a quick trip to the implementation of the UDMEM program. Try it both with and without `LOADHIGH`. You have DOS 5.0 and 6.0 to see the differences that `LOADHIGH` makes to its operation.

### Allocation Precautions

Each time `WINL2H` function `HE` loads a program for execution, it allocates memory for it as well as the UDMEM error file and under requests available RAM. For an EXE format file, the `blk` is a bit more specific, as is the amount of RAM needed. If this amount is not explicitly defined at link time, however, the EXE takes all available space.

It is important to remember that every byte of RAM each time they are loaded, whether they need it or not. This is particularly important to be sure that your programs turn themselves back to no more than the needed. There are going to be spawning other processes. Failure to do so will result in "out of memory" error messages. Low-level RAM and system controls.

Each time a program terminates normally, and returns control to its parent process, all RAM also allocated to it is freed. It is again becomes available for allocation. If the program terminates through a crash or `USR` errors, only part or possibly none of its memory is released to be used again.

It is important to remember that all available space while they are executing and can turn it back when they finish. Memory allocation for your programs can thus be handled automatically and in a way DOS uses. Unfortunately, getting one large block of memory at start up and having it deallocated for you at termination is often inconvenient because it means your program can't spawn



other processes. Modern C compilers usually include in their startup code the necessary bits to set their RAM usage (which is just what they require) to 64K, which never is larger, but at least one popular high-level language, Turbo Pascal, does not do this automatically. Instead, TP gives you a compiler option, "SM," that lets you specify how much memory to use.

Unfortunately, this option (in the DOS mode) details to "all available space." So in even condition results from attempts to VC or SPAWN a child process from C/P without using the SM option to make RAM available for the child process. Because of this, the Turbo Pascal developers probably gained a reputation for being frugal—actually, it was just not adequately documented.

In most high-level language programming, you won't use the three DOS RAM allocation functions directly but will use the `malloc` or `realloc` functions, which are really just wrappers that reflect you are using them indirectly.

The strategy in `malloc` is to obtain large blocks of RAM by using the DOS functions. Then divide it out into portions of much smaller portions as requested. Once RAM is allocated, the program never turns it loose. Elements allocated, even after calls to free, until the program returns to DOS. This strategy is the heritage of UNIX, where the allocation of system RAM was a time-consuming process. Under MS-DOS the reverse system. A number of improved performance packages, which replace the standard library versions of `malloc`, with more direct calls to the DOS functions, have appeared recently.

On the other hand, each block of memory allocated from DOS requires the additional 16-byte MBF, so all DOS allocations are consequently paragraph based. So if you want to allocate 4 bytes from DOS, you must use `malloc` to ask INE 21h function 48h for one paragraph, 16 bytes. In order to get smaller requests, DOS then actually needs two paragraphs, because you asked for plus one for the MBF, plus one for the paragraph. The point is that the simplest possible direct DOS memory allocation actually uses 32 bytes.

Please don't get confused from this discussion that `malloc` is normally a direct equivalent to the DOS functions. In fact, some programmers, including Borland's, erroneously report that the system is out of memory, you attempt to use the direct DOS functions, then subsequently attempt to use `malloc`—no more status returns. This failure happens because the compiler's library routines depend on unrestricted access to certain RAM. If you grab a block using DOS, the library routines cannot get more and so they report an error.

## RAM Allocation Strategies

When a program requests some paragraphs of memory and the memory manager has more than one block free, it's possible to satisfy that request in several different ways. These different ways of allocating memory are known as allocation strategies. DOS provides a function, INE 21h function 58h, to select different strategies.

This function hasn't always been documented. For example, it is described in Microsoft's MS-DOS Filesystem Programmer's Reference Manual for DOS 3.3. Only with publication of Microsoft's *MS-DOS Programmer's Reference* for DOS 5.0 did the subject come out of the closet.

Since INE 21h function 58h is only documented now, we won't spend nearly as much space on it as we did in some first edition of this book, when the allocation strategies capabilities could only be described as semi- or untested. At that time, the function permitted you to search only from first fit to "best fit" or "worst fit" strategies. With DOS 5.0, six more strategies were added. The original three operated on free or optional RAM; three of the new ones duplicate the original three, but search the 640K memory, loading into conventional RAM, and one that searches both the CMB and

Since this is a first edition, between the three original strategies and the six new ones to solve the question of which MBF call to traverse, I describe only the original three, because

**First-fit Strategy** The first fit strategy is the default action unless you explicitly change things. In the first edition of this book, I wrote that "Even if you do change strategies, DOS will change back to first fit whenever it loads a program, although it follows our selected strategy for all other loading actions." This statement is a simplification true for all versions of DOS. In most versions, as a matter of fact, the strategy that you select (rules of memory allocation) until you change it.

Another important point in my earlier description of the process was the statement that the search for a block of free space is sooner or later that way large enough was located. DOS actually maintains six the top-level block pointers in the MDA, one for the first block that can satisfy a request, one for the best fit choice, and one for the last block. DOS searches the entire MCB chain, or chains, if using the 20-bit E strategy for a MCB access, until reaching the end. Only then does it select which of the three pointers to return as the allocated block.

A major criticism for this seeming waste of effort. While it would be possible to stop short after finding a suitable first-fit strategy, were it not for doing so would prevent DOS from merging adjacent free blocks back into single larger blocks. Since each such a merge makes an additional 16 bytes of RAM (the MCB for the extra block) available for use, the cumulative effect can be significant. By going all the way to the end of the chain for each allocation, DOS makes sure that all free memory blocks are as large as possible.

The first fit strategy causes the memory manager to make the allocation from the first block that is large enough to satisfy the request. If the block is larger than requested, only enough is taken off the free list, the bytes used, and a new smaller block is created for the remainder.

In some circumstances DOS operation, there is usually only one such block in the system when a program is loaded. Because the loader often asks for all available RAM, no new block is created. Under these conditions, there is no difference between first fit and best fit strategies. If, however, the available RAM has become highly fragmented, and at the same time the block being allocated is small enough to fit in the first free block encountered, the first fit strategy uses that first block.

**Best-fit Strategy** The best fit strategy requires that the smallest block that will do the job be allocated, regardless of whether it is the first one encountered. As with the first fit strategy, the block is allocated from the free list, and any left-over space is put into a new, still free, block.

This approach guarantees that multiple allocations of small blocks don't fragment RAM unnecessarily. As long as blocks are released at approximately the same rate as they are allocated, the best fit strategy will be using the same small blocks over and over, leaving the larger blocks free to accommodate requests that require them.

As pointed out in the previous section, normal operation with only one or two blocks of RAM in a block is still effective, even in between first fit and best fit. If, however, you are programming atypical or that blocks use RAM management, that makes short-term use of large numbers of small blocks of RAM, you'll want to keep this strategy in mind. It could keep you from running out of RAM (except, really, by leaving some of the remaining free blocks is large enough to fill your latest request).

Having said this, it is important not to over-simplify it. In fact, as any textbook on operating systems will tell you, first fit is a most-always the correct strategy to use.

**Last-fit Strategy** Unlike either of the other strategies, the last fit technique is designed specifically for allocations that you expect to hang around for a long time—such as TSRs or device drivers.

When a block of RAM is allocated using the last fit strategy, the highest possible block of memory that can satisfy the request is assigned. Normally this is the highest part of the final block of free RAM. It is here that the memory is placed at the end of the MCB chain, so it never need to be searched if you wish to go back to the standard first fit strategy for subsequent normal allocations.

In the first edition, I claimed that "the last fit strategy is of limited usefulness," which moved reader Art Rothstein to respond with details of a significant use he had found for this strategy.

We use just it. We load a TSR that provides common services to various applications. Our applications span out either perhaps three levels deep. To save RAM, we link our applications to a stack of the Microsoft C++ MSVC++ heap. When there is sufficient room in the stack, MSVC++ attempts to expand the appropriate segment via INT 21h AHE4Ah Move of the code of a TSR provided by another vendor, as OIB and IIB files that are intended to be linked with our applications. We added just enough code to make the TSR itself provide calling interfaces from our applications. One of the functions provided by the TSR allocates memory via INT 21h A1148h. Suppose application X calls this function to spawn application Y. MSVC++ spawns the user heap to build a copy of the DOS environment to pass to DOS 4B11. If the TSR use first bit for its DOS 4B11 call is a local one would usually be right below application X's data segment preventing a virtual heap from expanding to satisfy the spawning. To avoid this problem, we change the allocation strategy to use it while the TSR has control.

Obviously, Microsoft knew what I was doing, so I chose to make all three strategies available. Since the details are now documented, I find the discussion of memory management here, and move on to process management.

## Process Management

The idea of a process is a separate executable program that may be loaded into memory, but that may or may not be executing currently. It is central to the operation of MS-DOS. The whole basis of TSR programming is that a process may be retained in residence after terminating, but TSRs are not the only processes, but DOS manages. Every program loaded for execution, including the command interpreter itself, is a process.

### Program Files and Processes

Before we get into the details of how processes are managed, let's first look at how program files are laid out and see how they relate to the eventual executing process.

#### The COM File Format

The original format for executable files inherited from the CP/M operating system was the COM file. This kind of executable file starts with all four segment registers containing the same value and execution always starts at the 0100h offset, the IP#, with the stack pointer set to 0001.

Because the file system is limited and advanced technical techniques can be developed, Microsoft has attempted over years to push this format into obsolescence. However, it closes to die. For small sizes of files, it remains an excellent choice. It is also useful for other special purposes, as we shall see toward the end of this chapter.

#### The EXE File Format

The EXE file format is like that of the COM file, allows multiple segments for both code and data. The format also supports relocation and simplifies the linking of object files into a single process. A means to develop systems that use the original EXE format, which is still the standard version used by DOS, provides a translation for the segmented executable format that is used in Windows and OS/2 programs to support dynamic linking.

While both the old and new formats have been officially documented in a number of places, we've found the source code updates for both sets of file headers to our UNDOCS-DOS.H file. Here is the appropriate excerpt:

```
typedef struct {
    WORD sig; /* ERE Program Header */
    WORD LeftOver; /* always 0x5A4D, 'M2' */
    WORD Pages; /* No. of bytes on last page */
    WORD Pages; /* No. of 512 byte pages */
};
```

```

WORD Items;           /* No. of RelocationTableItems */
WORD HdrSize;        /* in paragraphs */
WORD ReqSize;        /* in paragraphs */
WORD DesSize;        /* in paragraphs */
WORD InitSE;         /* in relative paragraphs */
WORD InitSP;         /* at entry */
WORD ChkSum;         /* at entry */
WORD InitIP;         /* at entry */
WORD InitCS;         /* in paragraphs */
WORD FirstRelItem;  /* Offset from beginning */
WORD OverlayNbr;    /* at entry */
WORD ReservedE 16 ;
JLONG NewExe,       /* offset from front of file */
J @EE, far * LPXEXE,

```

EXE files generated by Microsoft's linker always use a 512-byte page size (at least for them). EXE HdrSize is always equal to a multiple of 512. Those generated by other linking programs may establish different sizes, so EXE HdrSize may vary.


When DOS loads an EXE file into memory, it reads the header information into RAM at a temporary location, as described by the program. The loader then adjusts all addresses in the program as indicated in the relocation table, which follows the header, sets the stack segment and pointer from the JLONG EXE InitSS and EXE InitSP, pushes the values of EXE InitCS and EXE InitIP onto the stack in preparation for returning to the program, releases its temporary memory, and (if it was separately allocated) the loader sets DS and ES to the process segment address, and enters the program by executing the instruction J @InitCS and EXE InitIP pushed onto the stack.

If the file is in the new Windows format, it contains the full header of the original format at its front, and EXE NewExe contains the offset to the additional header within the file. If the file is not in the new format, EXE NewExe is 0.

### The PSP: How It Identifies a Process

Ever since I was using DOS in a way management, there was no way to avoid mentioning the Program Segment Number (PSN). Now we can examine this critical DOS data structure in more detail. The PSP, a 256-byte block located immediately preceding the actual process memory, is the key to process management in MS-DOS.

Figure 7-20. The PSP and Its Relation to the Process



The PSP contains the DOS state, the handles, etc., for its process; the segment address of the PSP itself provides a unique within-the-single-tasking environment for which DOS was designed (see Chapter 6 for a discussion of nonunique PSPs under Windows Enhanced mode), identified by which the process is located and managed. Thus the PSP segment address is also known as the Process Identifier or PID value.

**History, Purpose, and Use** The Program Segment Prefix came to MS-DOS as a way of Scale Computer Products' 86-DOS which, for compatibility, took the idea from Digital Research's 8086 CP/M operating system. As MS-DOS developed through the years, however, the PSP evolved far far more than its CP/M equivalent. The PSP now embodies many of the ideas provided by other operating systems, such as UNIX and MVS, by the stack frame of the process directory. By proper use of information kept in the PSP, a process can pass data to other processes that it spawns, or it can return information back to its parent process. At the same time, many fields of the PSP are assigned holders from the days of CP/M.

The primary purpose of the PSP is to contain the system information necessary to start, or to finish a specific process. This information includes, but is not limited to, the address of the register to which control should be passed when the process terminates, the use of handles in which the process identifies its files and devices, the address of the environment space belonging to the process, the identity of the process's parent process, and last but not least, any arguments passed directly to the process when it was invoked.

A second important reason to provide methods of accessing DOS functions without INT 21h (this was much more important in earlier times than it is today) with CP/M, the interface to DOS Basic Disk Operating System, the use of the INT 21h functions was by way of a subroutine call to address 0005h. Consequently, to provide the same functionality without DOS, the PSP of every process contains a table of pointers to the dispatch area of MS-DOS itself, which has interesting enough variations even since Version 2.0 that is self-documented (Version 6.0). It does not seem to have any way to be right-popped.

Main UNIX variants provide function equivalents through a direct call to the user's stack frame, as so with the `exec` function of UNIX (equivalent to MS-DOS 2.0's special call to INT 21h) was added to the PSP of other DOSs. At various times was indicated in Lind's MS-DOS 2.1 *Technical Reference* as "for UNIX compatibility."

Neither of these provisions is very used. Most programs today simply use INT 21h, or, if a program is old enough to be obsolete, is equivalent.

**(Usually) Unique Process Identifier** MS-DOS, at 8086, only one critical process because it uses the associated PSP as a starting point for much of its file-management activity. Yet MS-DOS can be used in multitasking before—multiple processes. A key to multitasking, even operating system without any real-time process-scheduling, is change the current process.

Through the use of the MS-DOS's `kernel` or `kernel` that references to a system call, the process identifier is a 16-bit value of a process ID number. PID (PID is a 16-bit value that uniquely identifies each process, or, inside the system, regardless of whether it is active. However, the documentation does not explain precisely what the PID is).

Each process, at 8086, is given a unique segment address of the PSP associated with that process. Within the system's memory organization upon which the design of MS-DOS was based, the value of the PSP segment address is unique. Obviously, a PSP can be located if any specific segment address, so the PID identifies the process without possible ambiguity.

Unfortunately, when the systems used its multitasking capability, under Windows Enhanced mode, the unique process ID number is used to deal with the structure of the Windows SYSTEM.INI file, `taskkill` settings, by `system` section, `UniqueDOSPSPs` and `PSP` section. As a result, it

task. In DOS, as in any virtual machine, PID differentiates processes. Thus, the combination of VM ID and PSP is a unique identifier.

Your tasking MS-DOS provided two undocumented functions and one documented one to store or retrieve the PIDs of the current process, thus activating one or another set of data stored in different PSPs. As we have already seen, undocumented became official with version 5. The current process is stored in the INT21h Function 80h. The current process can be queried either with INT21h Function 51h or with DOS function 62h, with the equivalent function 62h.

It is important to understand that none of the PSP functions actually deals with the PSP of the program that called it. As we saw in chapter 9, that is, the real deal moves a 16-bit value either to or from the current DOS word in the Swapable Data Area maintained by DOS. If the value in that word is nonzero, the Get PSP functions set the incorrect value with no indication of error.

It now appears to be a great deal of confusion on this point. For example, even Duwan's *Microsoft MS-DOS Programming* states that "Function 62h allows a program to conveniently recover the PSP address at any point during its execution, without having to save that program entry."

That is, the Get PSP functions always return its value that was last established in the SDP. So, the PSP of the suspended or current process in DOS (not necessarily the PSP of the call suspension) is the PSP saved from the ISR that has been activated by an interrupt. Get PSP returns the PSP of the suspended process, not the ISR. That is what makes the Get/Set PSP functions significant. This provides a means for switching between multiple tasks in MS-DOS. It is often said that DOS is a preemptive, preemptive phase, which means that only one process owns DOS at a given time.

As a result, the current process is switched, whether by your own multitasking code or by a ISR, as a group of operations is executed that the current PID also be switched if any I/O activity is to occur. Otherwise, the elements listed consist of the current foreground process, rather than our own, they will be affected.

The three get/set PSP functions are described in further detail in chapter 9 on ISRs and DOS multitasking.

**Undocumented Areas of the PSP** Less than one-third of the 256-byte area in the PSP has been documented. I estimate this section supplies information about the remaining parts. Not all of them, however, have ever been put to use.

The information in Figure 7-21 is taken from our `INTSR.DOS.H` header file, which provides the program and DOS semantics described in this chapter. For information on PSP fields owned by Windows and by Novell NetWare, see Chapters 3 and 4.

**Figure 7-21 PSP Details**

```
#define ENV_IDX_PSP(psp_seg) (*((WORD far *) RK_FP(psp_seg, 0x2c)))
#define PARENT_PSP(seg)  (*((WORD far *) RK_FP(psp_seg, 0x16)))
typedef struct {
    WORD sig;           /* 0000 always (9 20 (INT 20) */
    WORD nextgraf;     /* 0002 first unused segment */
    BYTE skip1;       /* 0004 filler to align next */
    BYTE cpmCall[5];  /* 0005 CPM-like service call */
    FP ISV22;         /* 000A documented ISR vectors */
    FP ISV23;         /* 000E " (saved at start) */
    FP ISV24;         /* 0012 " */
    WORD ParentID;    /* 0016 PSP of parent process */
    BYTE KTable[20];  /* 0018 indices into $FT */
    WORD EnvSeg;      /* 002C environment segment */
    FP SevStk;        /* 002E saved $S SP at INT21 */
    WORD Nhdls;       /* 0032 nbr of handles avail */
    FP HtblPtr;       /* 0034 ptr to handle table */
    FP ShareChn;      /* 0038 SHARE's closing chain */
    BYTE skip2;       /* 003C unknown */
}
```

```

BYTE TruNamFlg,      /* 003D APPEND's TrueName flag */
BYTE Skip3a[2],     /* 003E Used by Windows? */
WORD Version,       /* 0040 Major, minor vers level */
BYTE Skip3b[6],     /* 0042 Used by Windows? */
BYTE WlddApp,       /* 0048 Windows OldApp Flag */
BYTE Skip4[7],      /* 0049 unknown */
BYTE Disp[3],       /* 0050 Unix-like dispatcher */
BYTE Skip5[2],      /* 0053 unknown */
BYTE ExtFCB[7],     /* 0055 extended FCB1 area */
BYTE FCB1[16],      /* 005C documented FCB areas */
BYTE FCB2[20],      /* 006C */
BYTE Tailc,         /* 0080 'command tail' count */
BYTE Tail[127],     /* 0081 start actual data here */
} PSP, far * _PPSP,

```

The `ENV` (M:PSP:psp\_seg) and `PARENT` (p:psp\_seg) macros make it simple to obtain segment addresses for the associated environment block and for the process' parent process. When using either a far and/or M: or p: to compute a segment address of your program's PSP from the global variable `_psp`

If you're ready to point to your PSP by using `MK_FP(psp,0)`, you can then access any of the PSP elements as structure members. For instance, if you're in the pointer as `PSPptr`, the user's parent's segment address could be obtained as `PSPptr->Parent ID`.

## DOS Termination Address

Although the DOS interrupt service vectors used in the PSP are for interrupt usage, `INT 22h` and one of the top 256 keys is a way to hook a process at termination time. In effect, you cause the process to terminate. This provides a DOS exit capability. The basic vector is `INT 22h` (the `INT 22h` vector). From the DOS as the termination address. What's so undocumented is the fact that the address of a PSP's tail of the memory of a interrupt service segment is the one used with the process terminates.

To hook this vector and cause your own code to be executed when the process terminates and before `dos_terminate` is called, the `myprog.exe` just use the routines shown in Figure 7-22, with your own code executed in the `DoHook` routine. Set a hook with `EN` pointing to the PSP. `DoHook` will be called automatically at termination time.

**Figure 7-22. Process Termination Hook**

```

SetHook PROC
MOV     AX,[ES:000Ah]           ; save old offset
MOV     word ptr ES:OldVec,AX
MOV     AX,[ES:000Ch]          ; save old segment
MOV     word ptr CS:OldVec+2,AX
MOV     AX,offset DoHook
MOV     [ES:000Ah],AX          ; set in new vector
MOV     AX,CS
MOV     [ES:000Ch],AX
SetHook ENOP
DidVec  DD     0                ; place for old vector

DoHook PROC FAR
; whatever you need to do is coded here
JMP     [CS:OldVec]            ; then chain to original
DoHook ENOP

```

The termination address stored at `INT 22h` in the PSP is just the segment address to the `Exec` function called `INT 22h` (hex offset `4E00h`) that the parent used to invoke this process. Obviously, then,

when DOS transfers control to this address, it is ready to return to the parent. Doffhook is grabbing control instead, so when Doffhook is executed, all memory allocated to the terminating process has already been released, allowing the memory containing the Doffhook code. All files have been closed and the current SP has been set to that of the parent. In DOS 3.0 and later, the registers have been restored to the values they had when the parent performed the EXEC.

Basically, all that Doffhook should do is to install any special handlers that the program had a chance to set that will not be left pointing to addresses that are no longer valid. No other actions should be attempted, and by all means no file access or other I/O should be tried since the context control code, depending on the parent program, is or more complex on exit processing you are better off using various system events in C. Note, though, that the C routines are called only for normal process termination and are not to be used in case of fatal errors or other abnormal conditions.

For a better approach to exit handling, see the discussion of INT 21h AX=1122h Process Termination Hooks in Chapter 8.

### Other PSP Fields

The first full (undocumented) area of the PSP is the word at offset 0016h, which contains the PID of this process, parent process. If this process is the current command interpreter, its own PID appears here. A user can usually respond to a PID that can be terminated by the FMT command. Were it not for this, a user could track back through these pointers from one PSP to the parent PSP and thus locate the master command interpreter. However, not users can do so by tracing this way to locate the current shell, which may not be the master. A normal value, INT 21h, can be used to find the master copy of COMMAND.COM; more details appear in Chapter 10 on command interpreters.

Immediately following the FARENT pointer, at offset 0018h, is the 20-bit handle table. If 1, it shows that this PSP represents a file into the System File Tables maintained by DOS. The first five entries are normally set up by the kernel routines to provide handles for stdin, stdout, stderr, stderr, and fdopen. Note that the first three handles refer to the same System File Table; see Chapter 8 on the DOS file system and for detailed OS. All unused handles have the value 0xFF.

The next index is a field area in the local control at offset 0021h, which the DOS dispatch code uses to save SS and SP each time this process calls INT 21h; see step 6 in Figure 6-7. Saving the stack location in the PSP rather than in DOS' own data area makes multitasking possible by permitting DOS to continue other processes assuming each process saves its own last halted, that is, creating a process stack on entry. However, MS-DOS itself has not yet taken advantage of this capability.

Immediately after the FARENT pointer, at offset 0022h, is a 20-bit handle table, which permits you to refer to the handle table and thus make more than 20 file handles available to your process. A documented DOS function, INT 21h function 67h, exists to manipulate this area. An alternative to function 67h appears in FHANDLE, in Chapter 8.

The next two bytes of this register are the word SHHDS at offset 0052h, which defines the number of handles on a file, similar to process, attempting to open another file or device when this many handles are already in use (ignoring DOS' own). The following four bytes, HIBIPTR at offset 0034h, are a 16-bit pointer to the bottom of the handle table. By default, SHHDS is set to 20 and HIBIPTR to PSP+0018h, thus describing the handle table of the PSP.

The fourth word, at offset 0038h, always seems 0xFFFF in DOS versions prior to 3.3. Later DOS versions insert this amount to a previous PSP when SHARE is in use, creating a chain of PSPs used by SHARE's shared program, close, always opened by a given machine.

At the end of the PSP is another 30-bit of the PSP to maintain its true name flag for INT 21h function B\*1, IS=005, created. At Version 3.0, DOS began storing the major and minor version values at offset 409h. NEXTVER uses only the low 7 bits, applications about the DOS version number. WINDOAP (to Windows) uses the byte at offset 48h, setting the bottom bit when running old, that is, DOS, applications. Finally, offset 55h is the start of an extended FCB field.



### Spawning Child Processes

As is discussed in detail in Chapter 10, every program run under MS-DOS can be thought of as a child process. The only process that is not loaded as part of the boot action—that is, the loader that is read in from the boot sector—is the disk's child, the ROM BIOS bootstrap routine.

A child process is simply a process spawned by some other process, which is called the parent. Again, except for the bootstrap loader code that initially brings your system into action, every process in the system is a child of some other process.

The bootsup loader spawns only one child, the command interpreter specified by the SHELL line of CONFIG.SYS, or COMMAND.COM by default if no SHELL is specified. This process is what most users perceive to be DOS itself. Each time a program is invoked on the command line that requires spawning is a child of the command interpreter for execution.

If the spawned process is itself a memo or other type of shell routine, it may be able to spawn children of its own, which then execute and let their own parent know when they should exit. Should control ever return to the bootstrap loader, the result is the error message "Bad or missing command interpreter" and a forced system rebooting. However, DOS 6.01 does prompt for the full pathname to the command interpreter if it is unable to locate the specified one while performing the initial program load or bootstrap process.

### Locating Parent Processes

From time to time, a program needs to be able to trace its ancestry. It is not always possible. For example, a child program such as COMMAND.COM is always its own parent. For exactly the same reason, see Chapter 11, note the "I am stopped" bit there. However, if the process is running as the child of another process that is a command interpreter such as COMMAND.COM, the job of locating its ancestor's state is a historical task, and is complicated.

### Locating Ancestors

One way to locate the holder of the PSP, the PARENT word described previously, is to give a program the ability to trace its ancestry to the point of the closest command interpreter that spawned the program, to be sure that it is its own parent.

Thus a program that needs to locate its ancestor only needs its own PSP, except for a parent process ID, to use the process's own parent's PSP. The program counts upward the parent ID which PARENT will directly point to the PSP that contains a lower PSP's file, the next command interpreter program encountered in the trace.

### Use of this Capability

A simple program in C that uses this capability to trace its ancestry appears in Figure 7-23.

**Figure 7-23** Tracing Process Ancestry

```

/*
 * ROOTS.C (with apologies to Alex Haley)
 * Trace Your Ancestry!
 * Jim Kyle, 1990
 * revised August, 1992, jk to include UNDOCDOS.H
 */
#include <stdio.h>
#include <sys/types.h>          /* required to get _psp for RS */
#include <unistd.h>

WORD parent, self;

main ( void )
{ self = _psp;                  /* start with own PSP value */
  parent = PARENT( self );
  do
    ( printf( "PID = %d, PARENT = %04X\n", self, parent ),

```

```

    self = parent;
  }
  while ( ( parent = PARENT( self ) ) != self );
  return 0;
}

```

The program simply copies its own PID into the variable "self" and then uses it and the defined macro "PARENT" to construct the parent's PID in "parent".

It continues the program loops reporting the values of the variables "self" and "parent" at each level and then redefining them both, until the program reaches the level at which the two values match. This is the command interpreter. At this point, the program returns. It's most instructive, by the way, to run this program from some environment, rather than from the command line, because that guarantees at least one level of ancestry before the command interpreter is reached. For example, we can run ROOTS inside of DEBUG inside another copy of DEBUG:

```

D:\DOS2\CHAP7>debug \dos\debug.exe roots.exe
-9
-9
PID = 6C4E, PARENT = 672E
PID = 672E, PARENT = 615D
PID = 615D, PARENT = 6139

```

Here, 6111 is ROOTS, 672E and 615D are DEBUGs, and 6139 is the current command interpreter shell, which is actually a secondary shell invoked from my text editor. Naturally, we could use the excellent clipboard utility UDMEM, especially the function program (m\_psp) to find the ASCII names of these ancestors.

## Device Management

In addition to memory, the operating system must manage all devices connected to the CPU, such as the disk drives, the keyboard, and any displays. DOS manages, and issues requests to, device drivers, which in turn talk to lower-level interfaces such as the ROM BIOS Basic Input/Output System, or a device controller.

### Why Device Drivers Exist

Older operating systems, and even MS-DOS 1.x, included all hardware-dependent code necessary to deal with input and output as an integral part of the system itself. This made it necessary to bring out versions 1.1 and MS-DOS when IBM made available the 400k double-sided floppy disk drive, and made it impossible to use any kind of hard disk concurrently on a DOS 1.x system. Improvements were obviously in order.

A major part of the upgrade provided by MS-DOS 2.0 was the installable device driver capability. The idea, some appeared to originate at MIT with the MULTICS mainframe system, found its way into DOS as was of fact into UNIX, but the idea was significantly improved on its way into DOS. The original idea concentrated all hardware dependencies into small modules that were separate from the main mass of operating system code, but that still required the user to rebuild and re-link the operating system in order to change any drivers. This was true for both MULTICS and UNIX. In DOS, on the other hand, all you do to change a driver is modify CONFIG.SYS and reboot. As we show below, the end of this chapter, even that process can be made simpler, and a new driver can be installed from the command line, with absolutely no change to the main operating system code itself.

An installable device driver is a code package that forms a self-contained unit capable of initializing itself and through which all communication to and from a specific hardware device can be channeled. The format of the driver and its command interface is specified by the MS-DOS documentation.

Devices as far as DOS is concerned, come in two flavors: known as character and block. Character devices are those that can deal with a single character at a time, such as the CRT, the keyboard, the printer, or the serial port. Block devices are those that take disk drives, must simulate a block of data and transfer it at once. Drivers for the two types are distinguished by a single letter in the device header starting with a word and show the fact that some of the command functions make sense for only one of the two: `read` or `write`, `seek` or `status`, and require identical interfacing techniques.

By separating hardware dependencies into these modules, only new drivers, rather than a complete operating system upgrade, need to be developed when a new hardware device becomes available, the new device then immediately becomes usable with any older system that can accept the driver.

### Hardware-Dependent Details

In general, these aspects are intended to be highly device-specific and vary from one device to the next. These include actions required to initialize the device and prepare it for use, those required to send data out, and those required to receive data from it.

You might think that some devices need only two of these groups, because you don't usually send data to a keyboard or receive data from a printer. However, the keyboard does have to receive certain commands from the operating system to acknowledge that its output has been accepted, and similarly the system needs to read status conditions from the printer. These requirements are I/O devices, not just I/O devices.

Other details that are associated with specific hardware items, rather than with generic logical functions, include port addresses through which communication is achieved, the handshake protocols used to transfer data to and from the device, and the actual bit patterns transferred to convey data and status.

All of these hardware details, if clearly so, are concentrated within the single driver for each device. In general, the DOS drivers themselves are grouped into a small collection of logical functions as specified in the DOS documentation.

**Logically Required Functions** The DOS documentation specifies 17 logical functions, all of which must be supported or responded to by every device driver (regular or whether that driver itself makes sense to be used or not), using the case of `modecheck` for a CRT as a case. Comments provide additional details to help determine which I/O requirements a given device has.

Normal operation of a system is implemented with a simple table. The driver uses a table with a number as the index to a table of bit addresses; this table may point to the index address. If the specific `read` or `write` operation is driver-specific, the table is treated just as data in appropriate status code with no other action performed.

**Congruence of Files and Devices** One of the most useful results of the device driver idea is that MS-DOS can use files and devices in exactly the same way. However, it is not exactly obvious why you might expect this to be a good idea. At first glance, files and devices don't seem to have a lot to do with each other.

Let's try some thought experiments. For both files and devices, what really counts is the stream of data that your program would not even know whether such a stream comes from a file, from the keyboard, or from some other source, just that they'll be that much simpler to deal with when some new type of input device arrives on the scene.

This approach has the obvious, less-obvious advantage, when a computer is required to read data from a source that requires a completely different programming technique, data from a keyboard that was received from a file, otherwise identical data from a disk.

Using a device that is not a keyboard as an input device can be used through the drivers, and a way to increase the speed be obtained from the CRT. If you are programming a file

time video image display system with hooker control, you'll be forced to go directly to the video display controller with your output and to use BIOS routines to read the keyboard without waiting until the operator presses ENTER.

Legacy disk programs that run under MS-DOS are able to take full advantage of the power offered by the new idea. This is not a limitation inherent in the idea itself, but rather an artificial one imposed by the design of MS-DOS and its failure to anticipate all future needs. Or maybe it's limitation inherent in the idea of device independence. Windows device drivers deal with this problem by returning to the use of separate specifications for display drivers, keyboard drivers, and so on.

One interesting by-product of the file device congruence is that all your named devices can be accessed as files in any disk directory. This comes about because the BIOS routines that open both devices and files always search for devices first. If a device name is the same as the name of the file you want to open, the device will be opened instead. Because most of the procedures that determine whether a given file exists depend on trying to open the file and then detecting the error if it cannot be opened, these routines show that any device exists as a file in any directory you happen to test. Nevertheless, the device does not show up in the directory listing.

Files can be used to test for the existence of a directory itself, because if you try to open a device by attempting to use it as a file in a nonexistent directory, the directory error occurs before the device access attempt. That error, in turn, indicates that the directory itself cannot be accessed; if the directory can be accessed, the device can also always be accessed. The following batch file uses this aspect of DOS devices:

```
echo off
ren %dir%.bat
if exist %dir% goto exists
echo No such directory
goto done
exists
echo Directory exists
done
C:\>DIR %E%>%dir%>fooper
No such directory
C:\>DIR %K%>%dir%>%dir%\%k%
Directory exists
```

### Tracing the Driver Chain


In order to operate at all, MS-DOS must provide at least a minimal set of built-in device drivers. Yet to take full advantage of expansion, it's necessary to be able to insert new drivers at will and to have the power of replacing an existing driver with a new version.

To do for hardware what `chain.sys` does for BIOS routines, the drivers are organized as a singly linked chain, with a defined starting point that is always at the same place with any specific DOS version: the location of `IBIOS` in any version of the disk, however. Each driver in the chain includes as part of its structure a pointer to the next one, and the end of the chain is signified by the value `FFFFh` in the offset position of the hardware stack. Unlike the MIB chain, this is a track-linked list. Figure 7-24 shows how the chain of device drivers is organized in DOS.

The original device chain is prebuilt in the hidden system file `IO.SYS` in PC-DOS, `IBMDOS.COM` in IBM-DOS. If you add drivers to your system using the `DEVICE=` command in the `CONFIG.SYS` file, they are patched into the chain by the initialization portion of `IO.SYS` each time you boot your system.

Subsequent sections of this chapter describe the detailed organization of the device driver chain, tell how drivers are initialized during system boot-up, and then show you how to locate the start of the chain for any version of DOS and how to trace the driver chain and find out what is in your system.

Figure 7.24. The Device Driver Chain in DOS



**Organization of the Device Driver Chain** The device driver chain is a singly linked list structure with a defined starting point. The link itself is a far pointer (32-bit segment offset format) that forms the first two bytes of each device driver; the starting point is the driver for the NU device.

The NU device's characteristics are the bit bucket for both input and output; any output sent to NU simply vanishes without trace, and any attempt to read input from this device encounters a permanent EOF condition. To install a device with these characteristics is rarely useful; NU also serves as the anchor location for the driver chain.

As always, NU's link pointer holds the address of the supplied CON driver—the default console keyboard/CR/LF device. This driver is located near the front of the DOS memory area, which normally is above address 000000. The NU driver, however, is located near the front of the DOS category itself, within the SdosVars structure, which is at a much higher address.

Because the DOS handle-processing routines know where the NU driver is located, they can trace through the chain to obtain any requested driver.

As a ready reference, the DOS routines always go through the device chain looking for a match between the name of each character device and the requested name, whenever any attempt is made to open a device for input or output. Only when a match is found in the driver chain does DOS search the directory for a named file. This makes it impossible to either create or access a file that has the same name as a device. It might be possible to develop a form of security system based on this fact: first create the file, then installing a device with the same name and providing a secure method for changing the device's name during operation.

Note that the character devices have names that are used in the search; block devices are referred to as files, accessed only by name. During a name search, block devices are simply skipped. Because the first match to a name ends the search, an existing driver is replaced simply by inserting the requested driver into the chain where it will be encountered first and being sure that it has the same name.

**How Drivers Are Initialized** When you add new drivers using CONTRSYS, each driver is added to the front of the chain as it is encountered. DOS copies the link values from NU into the new driver's link field then puts the new driver's address into the NU link member.

block and character device drivers are added into the chain in the same way. The search always begins at the NUL driver, guaranteeing that any new drivers added will be found before the built-in ones.

Device driver patching that inserts each driver into the chain is not done (though) until the last step of the installation routine. First, the installation routine calls the driver's own internal initialization code. If a problem occurs, the installation is stopped with an advisory message. If the initialization completes without error, DOS checks the driver's attribute word to determine whether the driver is for a character device or a block device. If it is for a character device, it is added to the chain immediately.

If it is for a block device, DOS checks the number of units installed by the initialization routine's `UNIT` system call, that signals DOS not to install the driver even though no errors were reported. Otherwise, DOS uses the same code to assign it the next drive letter in sequence; then creates a `DRIVER` system block and the Appendix to the device and I/O in the DPR in from initialization routine's initialization program. Next, DOS builds a mount directory structure (see the Appendix) for the driver, which gives the other side to the device driver. Only after all these steps are passed completed does DOS patch the driver into the chain.

The device's specific entries at the part several sectors into a file and specify the initialization parameters for `DEVICE` files. However, when you do this you must be warned: not all devices share the same BIOS routines for application to block devices. The code for `CONSOLE` SYS uses a routine for the disk DPR for each such device immediately following the `UNIT` system call, so that the address returned to DOS by the driver when it initializes a file is the `DISK` BIOS when the driver is a hard disk unit. Thus, if you have more than one disk device driver in the system, each should return different break addresses, and these addresses should not be too close to any code that will be needed after DOS calls the driver's initialization routine.

Finally, you can add block device drivers in the same file, which is not prohibited by the system BIOS. However, you must be aware that the order you must be sure that all the character drivers appear in the file before any of the block drivers for the same reason.

Finally, a word of caution: to add a replacement driver, one has to avoid these possible problems. Sometimes, a driver can be necessary, do otherwise. When that's the case, be very careful and check the initialization sequence closely to be sure that any unit break address pointer is not wiping out driver code.

**Locating the Start of the Chain** The start of the device driver chain, the start of the `MC` & `CR` data structure in the `INT` file, can be determined using the undocumented `INT 21h` Function 5Bh, `Get List of Lists`. The NUL device driver header, the actual header, and a pointer to the first entry in the chain is maintained in the `LIST` of Lists.

For DOS 3.0, the NUL entry begins 14 bytes past the address returned in `ES:BX` by `INT 21h` Function 5Bh. With DOS 3.0, the offset is 28h, but with 3.1 that came down to 22h, and there it has remained.

The following code fragment shows how to use `ES:BX` with the address of the NUL driver for DOS 3.1 and up, for earlier versions, it is the constant 22h to the appropriate value.

```
mov ah, 52h ; get List of Lists
int 21h
add ebx, 22h ; NUL driver offset, DOS 3.1+
```

**Tracing It Through** Once you have located the start of the device driver chain, actual tracing through the devices. To get specific information of DOS during an `OPEN` function is simple. The only complication is to use the `FILE` system to track between character and block devices and to report block devices differently because they have no names.

The sample program shown in Figure 7-25, written for MASM version 5.1 but usable with other assemblers that support the simplified segmentation directives, shows how simple it is.

**Figure 7-25 Tracing the Device Chain**

```

; DEV.ASM--for DOS 3.1+
.model small
.stack

.data
blkdev db 'Block: ' ; block driver message
        db '0 units)'

.code
dev proc

        mov     ah,52h ; get List of Lists
        int     21h
        mov     ax,es ; segment to AX
        add     bx,22h ; driver offset, 3.1 and up
        mov     di,seg blkdev
        mov     dx,offset blkdev

dev1:   mov     ds,ax
        lea     si,[bx+10] ; step to name/units field
        test   byte ptr [bx+5], 80h ; check driver type
        jz     dev3 ; is BLOCK driver

        mov     cx,8 ; is CHAR driver
        lodsb ; so output its name
dev2:   IFDEF INT29
        int     29h ; gratuitous use of undoc 005
ELSE
        push   ds
        mov   di,si
        mov   ah,2 ; Character Output
        int  21h
        pop  ds
ENDIF
        loop  dev2 ; then go look for next one
        jmp  short dev4

dev3:   lodsb ; get number of units
        add   al,'0' ; assumes less than 10 units'
        push ds
        mov  ds,di
        mov  blkcnt,al ; set into message
        mov  ah,9
        int  21h
        pop  ds

IFDEF INT29
dev4:   mov  al,13 ; send CR and LF to CRT
        int 29h ; gratuitous use of undoc 005
        mov  al,10
        int 29h
ELSE
dev4:   push  dx
        mov  ah,2 ; Character Output
        mov  di,13 ; send CR and LF to CRT
        int 21h
        mov  d.,10
        int 21h
        pop  dx

```

```

ENDIF
    mov     si,ba                ; back up to front of driver
    lodsw                   ; get offset of next one
    xchg   ax,ba
    lodsw                   ; and then its segment
    cmp    bx,0ffffh          ; was this end of chain?
    jne    dev1               ; no, loop back
    mov    ax,4c00h           ; yes, return to DOS
    int   21h
dev    endp
end    dev

```

When DEV ASM is run on a MS-DOS 6.0 system, it produces the following list of drivers. The bottom 15 are those contained in the hidden file `IO.SYS`; the 5 unit block driver controls drives A, B, and C, and the other 10 are the standard DOS devices. The less than > reports result from the fact that some of the new Microsoft drivers allow 12 units, while DEV assumes that no more than six are present. The duplication of block drivers is an artifact of having SMARTDRV.EXE installed; the system had only four disk drives.

```

MUL
Block: < unit(s)
Block: 3 unit(s)
DBL55Y$
CON
$M$K$X$D
$M$K$X$D
Block: < unit(s)
CON
AUX
PRN
CLOCK$
Block: 3 unit(s)
CON1
.P11
.P12
.P13
CON2
CON3
CON4
ENH$K$X$D

```

The duplicate CON is `CON:VJ:ANSI.SYS` because it appears in the chain ahead of the standard CON driver; it is always used.

It is somewhat surprising that `+` assembled with a `-DIN129` tag DEV ASM makes gratuitous use of undocumented DOS IN1 29h. IN1 29h is the *test put file* interrupt called from DOS when sending characters to a device whose attribute word has bit four set. It is tempting to use IN1 29h here because it does simplify the code just below label `dev2`. However, Chapter 1 notes that there really are places you should use documented DOS instead of undocumented DOS even when it seems like more trouble. Performance optimization has program access to those places. Although this program absolutely demands use of undocumented IN1 21h Function 52h, there are several reasons why it should not use undocumented IN1 29h:

- Essentially the same functionality is available with IN1 21h Function 2, although it may be a trifle slower.
- IN1 29h output is not redirectable. Because this program displays block devices using INT 21h Function 2, which is redirectable, using IN1 29h elsewhere means that running `DEV > TMP TMP` ends up displaying character devices on the screen and block devices in the file. Pretty silly!



Thus, DIV provides a nice demonstration of when undocumented DOS is needed and when it isn't. I leave some discretion here. Don't use undocumented DOS if you don't need to. End of lecture.

## Loading Device Drivers from the DOS Command Line

To complete what you've learned about DOS resource management, let's create a program you can use to load device drivers from the DOS command line, without having to edit CONFIG.SYS and reboot.

Ever have an MS-DOS program that required the presence of a device driver, and you wished you had a way to install the driver from the command-line prompt, rather than having to edit your CONFIG.SYS file and reboot the system? Of course you can be thankful that it's so much easier to reboot MS-DOS than it is to edit the kernel, which is what must be done to add a device driver to a VxD. With DOS 4.0, I borrowed the idea of installable device drivers from UNIX, it's often forgotten that DOS 4.0 had improved on the installation of device drivers by replacing the loading of a new kernel with the simple editing of CONFIG.SYS.

Still, most of us occasionally wish we could use type a command line to load a device driver and be done with it, so many installable device drivers. Also, developers of device drivers often wish they had a way to debug the initialization phase of a device driver. This type of debugging usually requires a debug device driver that loads before your device driver, one requires hardware in a real simulation, but one could load device drivers after the normal CONFIG.SYS stage.

Well, we've come a long way from the loading of MS-DOS device drivers, not only possible, it's relatively simple to implement, and you know it's a slow undocumented DOS. We present such a program, DIV400, written as a combination of C and assembly language. The program that follows is not the same one that appeared in the first edition of this book, for the slight modification over you that date as of page 16, November, 1991 issue of *Dr. Dobbs Journal*. This is a much more extensively debugged, revised code, corrects several problems that made the original unusable for loading block devices under DOS 4.0 or 5.0, this version more reliably determines the correct drive letter to use when adding new block devices.

Many readers have passed that feedback that helped improve DIV400 for this edition. Some of the most critical feedback was from Dr. Wright, who raised a number of problems to our attention and suggested corrections to those, which I gratefully acknowledge here. Another who spotted many of the same bugs, also wrote that caused serious software crashes to Crippled Forces to Geoff, the new version was tested for variations in the same way that DOS itself does earlier than by using the installation of a new kernel status letters value. Others whose comments helped greatly were Norman Powell, the Warner, William J. Wronberg, and Jay Long, whose testing helped aware of some bugs not noticed before, not if it had of the reported problems. In addition to signs, although we've corrected some of the DIV400 to use the UNDO, DOS 4.0 header file so that you would not be so confused by the array offset always sprinkled throughout the original.

To use the program, I assume you have to do a type DIV400 followed by the name of the drive to be loaded in a command line, just as you would supply them in CONFIG.SYS. For example, instead of placing the following in CONFIG.SYS:

```
device=c:\dos\ansi.sys
```

you would simply type the following on the DOS command line:

```
C:\>devload dos\ansi.sys
```

There are several ways to verify that this worked. First, you can write ANSI strings to CON and see if they are properly interpreted as ANSI commands. For example, after a DEVLOD ANSI SYS, the following DOS command should produce a DOS prompt in reverse video:

```
C:\>prompt $e[7m$g$g$e[0m
```

If it maximizes accuracy, you can tell that the new driver has been installed by running DEV and inspecting its display of the device chain. You can see your new driver at the top of the list, right after NUL, a subclass of any identically named drivers loaded earlier.

```
C:\UNDOK\KYLE>dev
NUL
QEMMSBAS

C:\UNDOK\KYLE>dev\od 3dos\clock.sys
C:\UNDOK\KYLE>dev
NUL
CLOCKS
QEMMSBAS
..
```


DEVLOD loads both character device drivers, such as ANSI.SYS, and block device drivers (drivers that support one or more drive units, such as VDISK.SYS, whether located in SYS or EXE files).

## How DEVLOD Works


Here is the basic structure of the DEVLOD program:

```
startup code (CD.ASM)
main (DEVLOD.C)
  Move Loader
  movup (MOVUP.ASM)
  Load_Drvr
  INT 21h Function 4B05h (Load Overlay)
  Get_List
  INT 21h Function 52h (Get List of Lists)
  based on DOS version number
  get number of block devices
  get value of LASTDRIVE
  get Current Directory Structure (CDS) base
  get pointer to NUL device
  Init_Drvr
  call DD init routine
  build command packet
  call Strategy
  call Interrupt
  Get_Out
  if block device
  Put_Blk_Dev
  for each unit
  Next_Drive
  get next available drive letter
  find last existing DPB
  INT 21h Function 55h (Translate DPB -> OPB)
  poke CDS
  link into OPB chain
  Fix_DOS_Chain
  link into dev chain
  release environment space
  INT 21h Function 31h (TSR)
```

DEVLOD's first job is to shove itself out of the way to the top of memory. This lets it load the device driver as low as possible, reducing memory fragmentation. Figures 7-26 and 7-27 diagram the significant steps that DEVLOD takes while running.

**Figure 7-26 First Three Steps of DEVLOD Action**

Memory maps not to scale

**Figure 7-27 Final Stages of DEVLOD Operation**

Memory maps not to scale

DEVLOD loads kernel drivers into memory using the documented `IOCS` function for loading overlays (NA 211) or on 4B05b. An earlier version of DEVLOD read the driver into memory using `IOCS` (call to `open` to get and close the driver, but this made loading FAT drive types difficult, using the `EXOC` function instead. `IOCS` handles both SYS and FAT files properly. It might appear that the `SETVOLSTATE` function 4B05b should be called first to perform SEFIVEK processing, but any device driver on this version-specific should never be loaded with another version of

DOS 4, general drivers that do not depend on undocumented features of specific versions do not perform version tests, so using `SETVER` amounts, in essence, to tying down the loader's safety valve.

`DEVIO.D` then calls our good friend `undocumented INT 21h Function 52h` to retrieve the `drive of IAS_DRIVE` pointer to the DOS Common Directory Structure array and a pointer to the `Next_Letter`. The location of these variables within the `System` varies with the DOS version number.

`DEVIO.D` requires a pointer to the `NTI` device because, as we saw earlier in this chapter when discussing the `DEV` program, `NTI` acts as the interface to the DOS device chain. Since `DEVIO.D`'s main purpose is to add new devices into this chain, `DEVIO.D` must update this linked list.

The DOS version indicates operation under `MS-DOS 1.x` or in the `OS-2` compatibility box. `DEVIO.D` puts out an appropriate message. Otherwise, `DEVIO.D` creates a pointer to the name base of the `NTI` device and the right bytes at that location are compared to the constant `NTI` to verify that the device represents what the pointer is correct.

In changing over the Appendix to this book, the astute reader may have noticed an undocumented DOS constant, `INT_DEV_IOCTL` (220h), which returns a `DEVX` pointer to the header of the second block driver. `NTI` is first. Since DOS mixes together all device driver headers, this effectively gets a pointer to any DOS driver chain. So, in code, `INT 21h Function 52h` is used. The reason is that `INT 21h Function 52h` is undocumented. `INT 21h Function 1220h` was meant to be called only from a DOS extension, such as a network calibration, and also segment registers set to DOS's kernel segment. You should check those other variables from `System`, in case you are loading a block device which is not known until later in the code, called the driver's `INI` routine.

Once `DEVIO.D` has a piece of this information, it sends the driver a initialization packet. This packet is created by the function `Init_Drv`, which is a packet with the `INI` command as the `device` `System` routine, and that is the driver's `Init` or `open` routine. The object of these routines are not to be stored in device driver header. `DEVIO.D` creates a function pointer, using the offset to call `Init` or `open` at that `Address`. `DEVIO.D` merely informs what DOS does when it loads a device driver.

Of course, since `INI` has, in essence, nothing to do, but fail out. It is important to note that you do not set a link of the driver to the DOS driver chain, so `FILE` can call the driver `INI` routine. Of course, `INI` succeeds. `DEVIO.D` keeps the process with its true mission, which takes place, oddly enough, in the function `Get_Chain`.

It is only at this point that `DEVIO.D` knows whether to save a block or character device driver, so it is here that `DEVIO.D` takes special measures for block device drivers. By calling `Put_Blk_Dev`, for which format is provided by the driver, `DEVIO.D` uses the well-documented DOS `DPB` structure to find the `device` `DPB` through the function call `INT 21h Function 53h`. To update `DPB` to `DPB`, `DEVIO.D` uses the `CDN` structure of the `System` device, and links the new `DPB` into the `DPB` chain. These new `DPBs` are added after `CDN` is used. This is the old `Free`. The `DPB`, `DPB`, and `CDN` are explained in detail in Chapter 8, which DOS the system. The key point is that in `Put_Blk_Dev`, `DEVIO.D` takes information returned by a block driver's `INI` routine, and produces a new DOS drive. This piece of the program to determine the `device` change, it is a `device`, because DOS 4.0 added one byte to the `DPB`, making it too large for the space allowed by the original `DEVIO.D` code.<sup>1</sup> The result was total system lockup when not independent to `INIT`. `DEVIO.D` must install a block device. This bug has been squashed.

When installing a block device driver, `DEVIO.D` needs a drive letter to assign to the new driver. As Chapter 8 explains in great detail, the `CDN` is an undocumented array of structures, sometimes also called the `Drive_Info` table, which maintains the current state of each drive in the system. The array is a `device` `type`, where `type` equals `IAS_DRIVE`. `DEVIO.D` pokes the `CDN` in order to install a block device driver.

In section `Next_Drive`, is where `DEVIO.D` determines the drive letter to assign to a block device if there is an available drive letter. One technique for determining the next letter, #inlined out with `DEVIO.D`, is simply to read the `Number of Block Devices` field, `nblkds`, out of the `List of`

lists. However, this fails to take account of SUBSTed or network redirected drives. Therefore, we instead walk the CDN looking for the first free drive. In any case, DEVLOD updates the `abdkdr` field, if it successfully loads a block device.

If the driver `name` and `is` is for a character device, DEVLOD checks two of the only `is` attributes, `word` (a device with no buffer) to update the `CON` and `CLK` KS pointers (a one list of 1 MS). If the character device being installed has its `SYN` bit set, the `CON` pointer is updated to point to the new driver's `cons` pointer (used by DOS to check for CTRL-C and CTRL-BREAK keystrokes). Similarly, if `CLK` is set, the `CLK` KS pointer is changed to reflect the address of the new driver. DOS uses this pointer for all references to the clock device, rather than going through the overhead of searching for it on its own. The original version of DEVLOD failed to maintain these two pointers and, as a consequence, DEVLOD ANSI SYS would disable Control-C checking (thanks to Geoff C. Hopell for letting us know).

And, when installing a character device, DEVLOD searches all SET entries for any character entry with a driver of the same name. If DEVLOD finds any, it replaces the SET's pointers to the driver with pointers to the new driver being installed. This is essential for maintaining proper operation of the critical error handlers, as pointed out by Dan Winter in the July 1992 issue of *The Third Journal*, also after coming to me on the original DEVLOD.COM. Steve Davis' code was not completely compatible with DEVLOD. I rewrote it, but the idea remains the same as the one being suggested.

Whether loading a block or character driver, DEVLOD uses the `break` address—the first byte of the driver's address space that can safely be turned back to DOS for reuse—returned by the driver. For block devices, this `break` address has been increased to include the newly created `APPLY` Ctl. This converts the `break` address into a control paragraph to be installed.

DEVLOD scans the DOS SYMNTI routine's actions to determine whether driver initialization was successful. For character devices, the `break` address is compared to offset zero in the driver segment, if the `is` flag is set to indicate a block device. For block devices, the number of times returned in the command packet is checked for initialization. If the number is set to zero, it means no `is` flag is any check and no constraints are set (not used by the driver).

DEVLOD then links the device header into DOS's linked list of driver loaders. It walks through `cdp4` to record the driver's success or failure, just to save the content of the `SET` driver's `break` and then to copy it into the `break` of the new driver, and finally to store its `is` address. The new driver's `is` `SET` driver is considered. Note again that the DOS-linked list is not a general list; rather, you know that the driver's [INI] succeeded.

Finally, DEVLOD gives your memory by releasing its own memory. The resulting hole in RAM causes no harm, as it is likely to pop up before a fact any program's successfully loaded uses it as its own memory. If the size of the environment is not large, such DEVLOD might also consider the `downward` DOS ISR function (INI 21), function 31h, to exit without releasing the memory now occupied by the driver.

## DEVLOD.C

Before your new edition, this routine, however accomplished, all this in less than 2,000 bytes of executable code, some constraints should be mentioned.

Many constraints are imposed or simulated by implementing DEVLOD as a COM program using the time-sharing system. The way the program moves itself up in memory, because of a check, after when the COM loader chooses the need to address all memory, each segment register.

For character devices, programs which are accounting, it's necessary to know every address that the program uses, so you can't execute from this prelinked code, any part of the memory supplied with the computer. Fortunately, in most cases that's not a serious restriction, nearly everything can be handled without them. To do so, by language strings take care of the few things that can't be done in C itself.

Borland makes it easy to completely sever the link to the runtime libraries. They provide sample code showing how to do so. Microsoft also provides such a capability, but its documentation is quite cryptic. Any of the Borland compilers for DOS can be used. The first version of DEVLOD was created for Turbo C, which the current version was done with Borland C++ 3.0 using native C syntax.

The test program as presented requires either Borland or Turbo C with its register pseudo variables, getch() and \_emit() features. As explained in Chapter 2, register pseudo variables such as AX, ptr, dx, and so forth, do not directly read or load the CPU registers from C. Both getch() and \_emit() simply emit bytes into the code stream; neither are actually functions.

Figure 7-28 shows the main program (DEVLOD.C).

Figure 7-28. The DEVLOD.C Program

```

*****
*   DEVLOD.C - Jim Kyle - 08/20/90   *
*   Copyright 1990,1992 by Jim Kyle - All Rights Reserved *
*   (minor revisions by Andrew Schulman - 9/12/90) *
*   (major rewrite by Jim Kyle - July-Aug 1992) *
*   (minor change by Jim Kyle - August 1993) *
*   Dynamic loader for device drivers *
*   Requires Turbo C or BC++, see DEVLOD.MAK also. *
*****/

#include <stdio.h>
#include <stdlib.h>
#include <dos.h>

#include "undocdos.h" /* defines DOS internal structures */

#define GETFLAGS _emit__(0x9F)
#define FIXDS _emit__(0x76,0x1F) /* PUSH SS, POP DS */
#define PUSH_BP _emit__(0x55)
#define POP_BP _emit__(0x5D)
#define GO DOS_getInterrupt( 0x21 )

unsigned _stklen = 0x200,
unsigned _heaplen = 0;

LPPSP PSPptr, /* used to access fields of the PSP */
char FileName[65], /* filename global buffer */
char *drvarg, /* points to char after name in cmdline */
WORD movsize, /* number of bytes to be moved up */
void (*far * driver)(); /* used as pointer to call driver code */
_PDDVR drvptr, /* holds pointer to device driver */
_PDDVR nulptr, /* pointer to NULL device (chain start) */
FP nuldrv, /* additional driver pointers */
FP nstdrv,
BYTE far * nblkdra, /* points to block device count in List */
WORD lastdrive, /* value of LASTDRIVE in List of Lists */
BYTE far * CDSbase, /* base of Current Dir Structure */
int CDSsize, /* size of CDS element */
WORD nu_seg, /* hold parts of ListOfLists pointer */
WORD nu_ofs,
WORD nu_ofs,
_PLDT _D_ptr,
DCHNPKT CmdPkt,
int SFT_size, /* used by nex FixSFT, set by GetList */

extern unsigned _psp, /* established by startup code */
extern unsigned _heaptop, /* established by startup code */
extern BYTE _osmajor, /* established by startup code */
extern BYTE _osminor; /* established by startup code */
void _exit( int ), /* established by startup code */
void abort( void ); /* established by startup code */

```

```

void movup( LPPSP, FP, int ); /* in MOVUP.ASM */

int TestName( LPODVR, LPSIT ), /* in TESTNAME.ASM for FixSFT */
void ChgSFT( LPSIT, LPODVR ), /* in TESTNAME.ASM for FixSFT */

void copyptr( FPP src, FPP dst ) /* copy far pointer */
{ *dst = *src; }

void exit( int c ) /* called by startup code's sequence */
{ _exit(c); }

void Put_Msg ( char *msg ) /* replaces printf(), uses DOS only */
{ _AH = 2; /* doesn't need to be inside loop */
  while (*msg)
  { _DL = *msg++;
    GO_DOS,
  }
}

BOOL Get_Driver_Name ( void )
{ char *nameptr,
  int i, } cmdlinesz;

nameptr = (char *)B(PSPptr->tailc); /* set up to parse */
cmdlinesz = *nameptr++;
if (cmdlinesz < 1) /* if nothing there, return FALSE */
  return FALSE;
for (i=0; i<cmdlinesz && nameptr[i]<' ', i++) /* skip blanks */
  ;
drvarg = (char *)&nameptr[i]; /* save to put in $I
for ( i=0, i<cmdlinesz && nameptr[i]>' ', i++) /* copy name */
  FileName[i++] = nameptr[i];
FileName[i] = '\0'; /* name copied, but good time to
for ( i, i<cmdlinesz && nameptr[i]>' ', i++) /* make all UC */
  if ( nameptr[i] >= 'a' && nameptr[i] <= 'z' )
    nameptr[i] &= 0x5F; /* take out case bit */
return TRUE; /* and return TRUE to keep going */
}

void Err_Halt ( char *msg ) /* print message and abort */
{ Put_Msg ( msg ),
  Put_Msg ( "\r\n" ), /* send CR,LF */
  abort();
}

void Move_loader ( void ) /* vacate lower part of RAM */
{ WORD movsize, destseg;
  movsize = _heaptop - _psp; /* size of loader in paragraphs */
  destseg = PSPptr->Haltqref; /* end of memory */
  movup ( PSPptr, NK_FP( destseg - movsize, 0 ),
    movsize << 4); /* move and fix segregs */
}

void Load_Drvr ( void ) /* Load driver file into RAM */
{ WORD handle,
  struct {
    WORD LoadSeg,
    WORD RelocSeg,
  } ExecBlock,

  ExecBlock.LoadSeg = _psp + 0x10;
  ExecBlock.RelocSeg = _psp + 0x10;
  _DX = (WORD)&FileName[0], /* ds:dx point to filename */
  _BX = (WORD)&ExecBlock, /* es:bx point to ExecBlock */
}

```

```

_ES = _SS;                                /* which, being local, is SS */
_AH = 0x4B03;                              /* load overlay (COM, SYS, EXE) */
GO_DOS;                                    /* BS is okay on this call */
GETFLAGS;                                  /* check what happened */
if ( _AH & 1 )                              /* if carry flag set... */
    Err_Halt ( "Unable to load driver file." );
}

void Get_List ( void )                    /* set up pointers via List */
( _AH = 0x52;                               /* find DOS List of Lists */
  GO_DOS;
  nulseg = _ES;                             /* DOS data segment */
  LOfs = _BX;                               /* current drive table offset */
  LOfptr = (LPLOLIMK_FPI nulseg, LOfs );

switch ( _osma or )                       /* NUL adr varies with version */
{
  case 0
    Err_Halt ( "Drivers not used in DOS V1." );
  case 2
    nbkdrs = (FP)0L;
    lastdrive = (LOfptr->ver.v2.lastdrv);
    NULptr = (LPDDR1E)(LOfptr->ver.v2.nul);
    nulofs = (WORD)NULptr; /* just the offset part */
    SFT_size = 0x2B,
    & mak;
  case 3:
    if ( _osminor == D)
    {
      nbkdrs = (BYTE far *)(&(LOfptr->ver.v3D.bk_dev));
      lastdrive = (LOfptr->ver.v3D.lastdrv);
      NULptr = (LPDDR1E)(LOfptr->ver.v3D.nul);
      nulofs = (WORD)NULptr; /* just the offset part */
      SFT_size = 0x38,
    }
    else
    {
      nbkdrs = (BYTE far *)(&(LOfptr->ver.v31up.bk_dev));
      lastdrive = (LOfptr->ver.v31up.lastdrv);
      NULptr = (LPDDR1E)(LOfptr->ver.v31up.nul);
      nulofs = (WORD)NULptr; /* just the offset part */
      SFT_size = 0x35,
    }
  case 4
    CDSbase = (BYTE far *)(&(LOfptr->ver.v31up.cds));
    CDSsize = sizeof CDS ); /* defined for DOS3.1 struct */
    break;
  case 5
  case 6:
    nbkdrs = (BYTE far *)(&(LOfptr->ver.v31up.bk_dev));
    lastdrive = (LOfptr->ver.v31up.lastdrv);
    NULptr = (LPDDR1E)(LOfptr->ver.v31up.nul);
    nulofs = (WORD)NULptr; /* just the offset part */
    CDSbase = (BYTE far *)(&(LOfptr->ver.v31up.cds));
    CDSsize = sizeof CDS ) + 7; /* V4.5 7 bytes bigger */
    SFT_size = 0x3B,
    break;
  case 10
  case 20
    Err_Halt ( "OS2 DOS Box not supported." );
  default:
    Err_Halt ( "Unknown version of DOS!");
}
}

```



```

void Fix_BOS_Chain ( void ) /* patches driver into DOS chr */
( WORD i,
  nuldvr = NK_FPC nu_seg, nulofs+0x0A ), /* verify drvr */
  drvptr = "NUL"
  for ( i=0; i<8; ++i )
    if ( *(BYTE far *)nuldvr+i ) = *((BYTE far *)drvptr+i)
      Err_Malt ( "failed to find NUL driver." );

  nuldvr = NK_FPC ( nu_seg, nulofs ), /* point to NUL driver */
  drvptr = NK_FPC ( _psp+0x10, 0 ), /* new driver's address */

  copyptr( (FP)nuldvr, (FP)&nuldvr ), /* hold old head now */
  copyptr( (FP)&drvptr, (FP)nuldvr ), /* put new after NUL */
  copyptr( (FP)&nuldvr, (FP)drvptr ); /* and old after new */
}

// returns number of next free drive, -1 if none available
int Next_Drive ( void )
{
#ifdef USE_BLKDEV
  return (nblkdrs && (nblkdrs < lastdrive)) ? nblkdrs -1,
#else
  /* This approach takes account of SUBSTed and network-redirector
   * drives by finding the first unused entry in the CDS structure
   */
  LPCDS cds;
  int i=0;
  for ( cds=(LPCDS)CDSbase, /* start at front of struct */
        i<lastdrive; /* go all way through it */
        ++, ((BYTE far *)cds)+CDSsize) /* count up */
    if ( ! cds->flags ) /* found a free drive */
      break;
  return ( i == lastdrive ) ? -1 : i, /* return number, or -1 */
#endif
}

/* This routine initializes the device driver and returns TRUE
 * if all went well. If initialization fails for any reason,
 * this function returns FALSE and the driver will not be linked
 * into the chain maintained by DOS.
 */
BOOL Init_Drvr ( void )
{ WORD tap;
  drvptr = NK_FPC ( _psp+0x10, 0 ), /* new driver's address */
  CmdPkt command = DD_INIT, /* defined in UNDOCDOS.H file */
  CmdPkt hdrLen = sizeof( BDCMDPKT );
  CmdPkt.unit = 0;
  CmdPkt.status = 0, /* clear status just in case */
  CmdPkt.inputs = (WORD)drvarg, /* points into command line */
  CmdPkt.inoscg = _psp,
  CmdPkt.NextDrv = (BYTE)Next_Drive(); /* for block dev init */
  if( CmdPkt.NextDrv == 0xFF &&
      (drvptr->attr & CHAR_DEV) == 0 )
    { Put_Msg( "Current Directory Structure is full, cannot install.",
              return FALSE;
    }
}

if ( _osmajor >= 5 )
  ( /* In DOS 5+, DOS passes the device driver irEndAddress
    (see DOS Programmer's Reference, p. 400) */
    CmdPkt.brkofs = 0;
    CmdPkt.brkseg = 0x4000, /* allow all RAM */
  )
}

```

```

tmp = drvptr > stratofs, /* STRATEGY pointer in driver */
driver = MK_FP( FP_SEG( drvptr ), tmp ),
_ES = FP_SEG( void far *) & CmdPkt ),
_BOX = FP_OFF( void far *) & CmdPkt );
(*driver)(); /* set up the packet address */

tmp = drvptr > introfs, /* INTERRUPT pointer in driver */
driver = MK_FP( FP_SEG( drvptr ), tmp );
(*driver)(); /* do the initialization */

/*
 * In the first edition version of DEVL0D, this function checked
 * the status code in the command packet to determine whether the
 * installation had failed. Actually, the status code is NOT
 * checked by DOS itself. Thanks to Geoff Chappell for pointing
 * out that SFSINIT does not check the status returned by drivers
 * after initialization. A block device driver is not returned if
 * its unit count is found to be zero, while a character device
 * should set its break address to offset 0 in its load segment
 */
return( drvptr > attr & CHAR_DEV ?
MK_FP( CmdPkt.brkseg, CmdPkt.brkofs) != MK_FP( FP_SEG( drvptr ), 0) :
CmdPkt.nunits != 0 );
}

```

\* This routine looks far more complicated than it actually is.
 \* It's used only when block device drivers are being installed,
 \* and does the housekeeping of DPBs and CDS entries that such
 \* device require. Major changes were made here in the second
 \* edition of the program, to accommodate changes in DPB size that
 \* happened at Version 4.0 but went unnoticed until DOS5 appeared.
 \* Special thanks are due Don B. Wright, Nathaniel Polish, and
 \* Geoff Chappell for spotting problem areas here.
 \*
 \* This routine returns FALSE if all goes well, or TRUE if any
 \* error condition is detected.
 \*
 \*

```

BOOL Put_Blk_Dev ( void )
{
  int newdrv,
  int i;
  int retval = TRUE; /* pre-set for failure */
  int BufferSize,
  int unit = 0;
  _DPB o_dpb, new_dpb, endmark (LPDPB)0xFFFFFFFFL,
  _PCDS cds,
  _PDWR new_driver (LPDWR)MK_FP( _psp+0x10, 0 );

  if ((Next_Drive() == -1) || CmdPkt.nunits == 0)
    return retval; /* cannot install block driver */
  /* (CmdPkt.brkofs != 0) /* align to next paragraph */
  {
    CmdPkt.brkseg += (CmdPkt.brkofs >> 4) + 1;
    CmdPkt.brkofs = 0;
  }
  while (CmdPkt.nunits - i) /* repeat this loop for each unit */
  {
    if ((newdrv = Next_Drive()) == -1)
      return TRUE; /* no room for another drive, quit */
    if( !nbldrs ) /* if not a null pointer, */
      (*nbldrs)++; /* ...tally into drive counter */
  }
}

```

#ifdef ORIGINAL

```

/* Tell DOS to get the DPB of the last drive in CDS. This
 * technique, used in the first version of DEVL0D, creates a
 * problem if the final drive is a JOINed or SUBStituted entry in
 * the list. The alternate method of finding the last DPB is

```

```

* not subject to this problem, but may be a bit slower. The
* address of the last drive DPB is saved in oldDPB"
*/
_AH = 0x32; /* get DPB of last drive in CBS */
_DL = newdrv,
GO_DOS;
_AX = _DS, /* save segment to make the pointer */
FIXDS,
oldDPB = MK_FP(_AX, _BX), /* this is base address of DPB */
#else
/* Trace out entire chain each time around the loop. While
* this is possibly slower than the original method, it will
* not be deceived by CBS entries. The address of the last
* DPB is saved in 'oldDPB'. Note that only the offset words
* of the link pointers are compared, DOS itself does not put
* "endmark" in the segment word, although this routine does
*/
oldDPB = LOLptr->dpb, /* always start at first DPB */
if (_osmajor < 4) /* trace through to the end */
while( (WORD)oldDPB->ver.v3.next != (WORD)endmark )
oldDPB = oldDPB->ver.v3.next,
else
while( (WORD)oldDPB->ver.v45.next != (WORD)endmark )
oldDPB = oldDPB->ver.v45.next,
#endif

/* Tell DOS to create the DPB, passing it BPB info from a
* list of near pointers passed back by the driver. Note
* that DS must be set after all memory references are done,
* because it's used to access the globals. Similarly, AX
* must be set after all segment registers, because it is
* used to load the segreg
*/
newDPB = (LPDPB)MK_FP(CmdPkt.brkseg, 0),
_SI = *WORD far *MK_FP(CmdPkt.inpseg, CmdPkt.inpofs);
_ES = CmdPkt.brkseg, /* ES:BP is DPB address to use */
_DS = CmdPkt.inpseg; /* DS:SI is adr of BPB to read */
PUSH_BP, /* save stack frame pointer */
_BP = 0; /* BPB offset value */
_AH = 0x53, /* build the DPB for this unit */
GO_DOS,
POP_BP; /* restore stack-frame pointer */
FIXDS,

/* Check to be sure that block sector size is acceptable
* If bigger than BUFFERS were built for, refuse to install.
* the driver. No such check was made in the original
* version of this program. Geoff Chappell spotted the omission.
*/
switch( _osmajor ) /* get BUFFERS size from LOL */
{
/* location in LOL will vary... */
case 2:
BufferSize = LOLptr->ver.v2.secsiz;
break;
case 3:
if ( _osminor == 0 )
BufferSize = LOLptr->ver.v30.secsiz;
else
BufferSize = LOLptr->ver.v31up.secsiz;
break;
case 4:
BufferSize = LOLptr->ver.v31up.secsiz;
break;
case 5:
BufferSize = LOLptr->ver.v31up.secsiz;
break;
default:

```

```

Err_Malt ( "Unknown version of DOS!"),
)
if( newDPB->bytes_per_sect > BufferSize )
    return TRUE; /* got out FY too big */

/* Now set the DPB address into the old last DPB's link
 * address field ("next" pointer).
 */
if ( _osmajor < 4 ) /* link new DPB to chain */
    oldDPB->ver.v5.next = newDPB,
else
    oldDPB->ver.v45.next = newDPB;

/* Set up the Current Directory Structure for this drive and
 * tag it as a physical drive (Clear IFS area if DOS V4 or
 * higher
 */
if ( _osmajor > 2 ) /* Version 2 did not use CDS */
{
    cds = (LPCDS)(CDSbase + (newdrv * CDSsize));
    cds->flags = CDS_PHYS, /* defined in UNDOCDS.H file */
    cds->ddb = newDPB, /* set DPB adr into CDS */
    cds->in_loc.start_cluster = 0xFFFF, /* not accessed yet */
    cds->in_loc.ffff = -1L,
    cds->slash_offset = 2, /* start in root directory */
    if ( _osmajor > 3 ) /* zero out IFS stuff */
    {
        *(WORD far *)&(cds->slash_offset)+1 = 0,
        *(WORD far *)&(cds->slash_offset)+2 = 0,
        *(WORD far *)&(cds->slash_offset)+3 = 0,
        *(BYTE far *)&(cds->slash_offset)+8 = 0,
    }
}

/* Finally, set up pointers for the DPB and the driver so
 * that they can find each other, and adjust space
 * reservations so that the BPB won't be wiped out upon
 * return to DOS. Step the BPB list pointer in case the
 * driver has multiple units.
 */
newDPB->drive = newdrv, /* set in drive number */
newDPB->unit = unit++, /* and also the unit number */
if ( _osmajor < 4 ) /* Versions 2 and 3 are alike */
{
    newDPB->ver.v3.driver = newdriver,
    newDPB->ver.v3.next = endmark,
    EndPkt brkseg += 2, /* was 32 bytes, exact fit */
}
else /* Versions 4 and 5 are alike */
{
    newDPB->ver.v45.driver = newdriver,
    newDPB->ver.v45.next = endmark,
    EndPkt brkseg += 3; /* 33 bytes each now */
}
EndPkt inpofs += 2, /* point to next BPB pointer */
/* end of units loop */
return FALSE; /* all went okay */
}

/* This function is called for a character driver only. It searches
 * every entry in the SFI, looking for any reference to the named
 * driver, and if such a reference is found, modifies 6 bytes in the
 * SFI entry to point to the just-added new driver. This is necessary
 * for proper operation of the critical error handler. It is based on
 * code by Dan Winter, published in the July, 1992 issue of Dr. Dobbs's
 * Journal, but Dan's original code was not compatible with DEVLDB.
 *
 * SFI size is established by the GetList() function earlier.
 */

```

```

void FixSFT ( void )
{ LPSFTB blk, next_blk,          /* use typedefs for simplicity */
  LPSFT sft;
  WORD num_items;

  for( next_blk = LOLptr->sft;    /* first SFT block */
        (WORD)next_blk != 0xffff, ) /* check all blocks */
  { blk = next_blk;              /* start this block */
    next_blk = blk->next;        /* link for next one */
    num_items = blk->here;       /* get size of block */
    sft = B[blk->first];         /* first SFT in it now */
    while( num_items-- > 0 )     /* test all items here */
    { if( sft->ver.v2.numera )    /* is this SFT in use? */
      { if( TestName( drvptr, sft ) ) /* yes, match? */
        ChgSFT( sft, drvptr ); /* yes, fix the pointer */
        (char far *)sft += SFT_size; /* to next item in list */
      }
    }
  }
}

/* This function is called only when the driver has been fully
 * installed with no detected errors. If the driver is a block
 * device, Put_Blk_Dev() is called to create its drive letter
 * and CDS and DPB structures, otherwise the COM and CLOCK
 * pointers in the List of Lists are updated as applicable. If
 * the block device installation fails, DEVLOD quits without
 * linking the driver into the DOS chain. Otherwise this driver
 * is put at the head of the chain right after MUL, and the program
 * returns to DOS, leaving the driver resident. In either case,
 * this function will never return to the main() procedure
 */
void Get_Out ( void )
{ WORD temp;

  temp = drvptr->attr;           /* attribute word */
  if( (temp & CHAR_DEV) == 0 ) /* if block device, set up this */
  { if( Put_Blk_Dev() )         /* fails if cannot do so */
    Err_Halt( "Could not install block device" );
  }
  else                           /* not block, check for updates */
  { if( (temp & IS_STDBIN) )
    LOLptr->con = drvptr;        /* this is for CTRL C checking */
    else if( (temp & IS_CLOCK) )
    LOLptr->clock = drvptr;      /* this is for fast time access */
    FixSFT();                   /* Ben Winter's fix for SFT */
  }

  Fix_DOS_Chain(),              /* all okay so patch into DOS */

  _ES = PSPptr->EnvSeg,          /* release environment space */
  _AH = 0x49;
  GO_DOS,
  PSPptr->EnvSeg = 0,           /* zero out the address in PSP */

  /* then set up regs for KEEP function, and go resident */
  temp = (cmdPkt brkofs + 15), /* normalize the offset */
  temp >>= 4,
  temp += cmdPkt brkseg,       /* add the segment address */
  temp -= _psp;                /* convert to paragraph count */
  _DX = (WORD)temp;            /* paragraphs to retain */
  _AX = 0x3100;                /* KEEP function of DOS */
  GO_DOS,                       /* won't come back from here! */
}

void main( void )               /* usual argc, argv not used! */

```

```
( PSPptr = (LPPSP)MK_FP( _psp, 0 ); /* create global ptr */
if( !Get_Driver_Name() )
    Err_Halt( "device driver name required." );
Move_Loader(), /* move code high and jump */
Load_Drvr(), /* bring driver into freed RAM */
Get_Inst(), /* get DOS internal variables */
if(Init_Drvr()) /* let driver do its thing */
    Get_Out(), /* check init status, go TSR */
else
    Err_Halt( "Driver initialization failed." );
)
```

### MOVUP.ASM

The small assembly language module MOVUP (Figure 7-29) contains only one function, movup. Recall that in order not to fragment memory, DEVLDR moves itself up above the area into which it loads the driver. The program accomplishes this feat with movup.

Figure 7-29. MOVUP.ASM

```

-----[
NAME      movup
]-----[
MOVUP.ASM -- helper code for DEVLOD.C
]
[
Copyright 1990 by Jim Kyle - All Rights Reserved
]-----[
]-----[

_TEXT SEGMENT BYTE PUBLIC 'CODE'
_TEXT ENDS

_DATA SEGMENT WORD PUBLIC 'DATA'
_DATA ENDS

_BSS  SEGMENT WORD PUBLIC 'BSS'
_BSS  ENDS

GROUP _TEXT, _DATA, _BSS

ASSUME CS:_TEXT, DS:GROUP

_TEXT SEGMENT BYTE PUBLIC 'CODE'

;-----;
;
; movup( src, dst, nbytes )
; src and dst are far pointers, area overlap is NOT okay
;-----;

PUBLIC _movup

movup PROC NEAR
    push    bp
    mov     bp, sp
    push    si
    push    di
    lds     si, [bp+4]          ; source
    les     di, [bp+8]          ; destination
    mov     bx, es              ; save dest segment
    mov     cx, [bp+12]         ; byte count
    cld
    rep     movsb              ; move everything to high ram
    mov     ss, bx              ; fix stack segment ASAP
    mov     ds, bx              ; adjust DS too
    pop     di
    pop     si
    mov     sp, bp
    pop     bp
    pop     dx                  ; Get return address
movup ENDP

```

```

push    bp                ; Put segment up first
push    dx                ; Now a far address on stack
_movup  ENDP
_TEXT  ENDS
end

```

### TESTNAME.ASM

New in this modified version is the assembly file `TESTNAME.ASM` (Figure 7-30), which provides two small routines that greatly simplify the ESNL patching procedure, described by Dan Winter in his July, 1997 letter to *The Linux Journal*. `TestName` takes far pointers to the new driver and to the current SFT entry; it then compares the 8-byte name fields, returning one if all eight bytes match and zero otherwise. `ChgSFT` takes the same two far pointers and corrects the affected fields of the SFT entry to reflect the new driver's address.

Figure 7-30. TESTNAME.ASM

```

;-----
; int TestName( LPDDVR, LPSFT )
; src and dst are far pointers, TRUE if 8-byte match found
;-----
PUBLIC  _TestName

_TestName PROC    NEAR
push    bp                ; save regs
mov     bp,sp
push    si
push    di
push    ds
lds     si,[bp+4]         ; get pointer to driver header
add     si,10             ; offset to name field
les     di,[bp+8]         ; pointer to SFT start
add     di,32             ; offset to name field
mov     cx,8              ; name length
cld
repeq  cmpsb             ; go while equal
mov     al,[si-01]        ; last chars tested
sub     al,[di-01]
add     al,0FFh          ; CY if nonzero
sbb    al,al             ; 0 or FF
inc     al                ; 1 if match, else 0
pop     ds
pop     di
pop     si
pop     bp
ret

_TestName ENDP

;-----
; void ChgSFT( LPSFT, LPDDVR )
; modifies SFT to point to new driver
;-----
PUBLIC  _ChgSFT

_ChgSFT PROC    NEAR
push    bp                ; save regs
mov     bp,sp
push    di
les     di,[bp+4]         ; get SFT address into ES:DI
add     di,7              ; offset to v.type.devdrv field
mov     ax,[bp+8]         ; offset of LPDDVR

```

```

mov     dx,[bp+10]      , segment of LPDDR
cld
stosw  , store offset
xchg   ax,dx
stosw  , store segment
xchg   ax,dx
stosw  , store offset
pop    di
pop    bp
ret

```

\_ChgST1 ENDP

We're both of these functions normally would be more easily done using C library functions such as `memcpy()`. If DEVLOD's used from far pointers with linear procedures made it much easier to perform the tasks essential. In any case, by building the field objects into this code, I was able to do things I could not do easily with combining casts, although at least one remains in the base 10 conversion of DEVLOD's version; the size of the SEI depends upon the DOS version, making it necessary to do a 4-byte arithmetic by casting the SEI pointer to `char *` before adding the table size.

## CO.ASM

The `CO.ASM` file is a CO-ASM file that has been extensively modified from start-up code provided with the DOS 3.31. It is an assembly code form part of every C program and provides the link between the DOS operating system and the program itself. Normal start-up code, however, does not move into the stack's above version. This code generates argument list sets up pointers to the stack, and arranges things so that library functions can operate. It also arranges for automatic freeing of the stack. These operations would DEVLOD's non-rotate.

Since the program has no need for any of these actions, our CO-ASM module omits them. What's left is the DOS 3.31 DOS error case and runs the RAM's used by the program down to the memory limit, and the memory sub-system. It sets the returned value onto the stack and calls `exit()`. A `ret` is used to return the success code, which is derived from the return value from main.

The `main()` function is performed; this module establishes global variables, assigns operators, and sets up each DEVLOD's copy or operation. Values for these variables are required from DOS's 3.31's unsegmented capabilities, and stored in the variables before control reaches `main()`.

The CO-ASM file is concerned with establishing an environment with which DEVLOD can run. It is with the actual objectives of DEVLOD itself, which is limited to using pointers. Both the source and the OBJ file are included on the companion diskette, however.

## Make File

Since this sample program includes two assembly language modules, in addition to the C source, a `MAKEFILE` file greatly simplifies invocation. Figure 7-31 shows one for use with Borland's `MAKE` utility.

Figure 7-31 The DEVLOD Makefile

```

# makefile for DEVLOD.COM created 05/25/90 - jk
#                                     last revised 08/20/92 - jk
# can substitute other assemblers for TASM, TCC for BCC

AS = D:\BC\BIN\TASM
CC = D:\BC\BIN\BCC
LL = D:\BC\BIN\LINK

devlod.com : devlod.obj c0.obj movup.obj testname.obj
$(&CC) c0 movup testname devlod /t/m/t,devlod.com

c0.obj : c0.asm

```



```
$(AS) c0 /t/mx/la;
```

```
movup.obj:    movup.asm
$(AS) movup /t/mx/la;
```

```
testname.obj: testname.asm
$(AS) testname /t/mx/la;
```

```
devlod.obj:   devlod.c
$(CC) c ms -P- devlod.c
```

You do, of course, need to change the three macro definitions to reflect the drive and path for your own compiler installation. Then you can simply type "MAKE DEVLOD MAK". The `r` option switch included in the LINK command line ensures that the linker generates a COM file, rather than the more usual EXE format.

### How Well Does DEVLOD Work?

A fitting conclusion to this chapter is to use some of the utilities developed earlier: UDMEM and DEVX to see what my system looks like after I've loaded up a couple of device drivers with DEVLOD. The report is shown in Figure 7.32.

### Figure 7.32. Testing DEVLOD

```
D:\VDOS2\CHAP7> devlod c:\ramdrive.sys 512 /e
```

```
Microsoft RAMDrive version 3.06 virtual disk E:
Disk size: 512k
Sector size: 512 bytes
Allocation unit: 1 sectors
Directory entries: 04
```

```
D:\VDOS2\CHAP7> devlod c:\ansi-uv.sys
```

```
D:\VDOS2\CHAP7> udmem
Seg  Owner  Size
0253 0008 092A ( 37536)  DOS Data Segment
      Seg Size Type
      ----
      0254 0041 Device Driver (386MAX)
      0296 0015 Device Driver (386LOAD)
      02AC 0747 Device Driver (STORDRV) [26 FA FE ]
      0PEE 0015 Device Driver (386LOAD)
      0A04 0050 System File Tables
      0A62 0005 FCBS
      0A68 0020 Buffers
      0AB9 0037 CDS Table
      0AC1 008C Stacks [02 0A 0B 0C 0D 0E 70 72 73 74 76]
007E 0008 0004 ( 64)  DOS Code area
0083 0084 0010 ( 256)  Env at 0314
0094 0080 0005 ( 80)  Free
009A 0890 0147 ( 5232)  No Env Segment /f:2048 /l 20
0CE2 0CEF 0008 ( 176)
0CEE 0CEF 0010 ( 256)  Env at 0A36
0CFF 000F 0008 ( 176)
0D00 0080 0001 ( 16)  Free
0D0D 000E 005A 1440)  No Env Segment c:\ramdrive.sys 512 /e
0D68 0069 0075 1872)  No Env Segment c:\ansi-uv.sys [1B 29 2F ]
0DDE 0007 1268 ( 75392)  Env at 0D00 D \VDOS2\CHAP7\UDMEM.EXE [00 05 17 ff ]
2D47 0000 7FB7 (523120)  Free [30 E6 E9 EC EF FA F5 F9]
9FFF
UMB Chain
CB00 FFFF 05BE ( 22752)  386LOADED Driver [13 15 28 ]
CDBF FFFA 0004 ( 64)  386MAX UMB control block
```

```

C094 FFFE 0205 ( 8272) 386MAX UMB
CF9A FFFA 0004 ( 64) 386MAX UMB control block
CF9F FFFE 0022 ( 544) 386MAX UMB
CFC2 FFFA 0004 ( 64) 386MAX UMB control block
CFCT FFFE 0022 ( 544) 386MAX UMB
CFEA FFFA 0004 ( 64) 386MAX UMB control block
CFEF FFFE 0042 ( 1056) 386MAX UMB
D032 FFFA 0004 ( 64) 386MAX UMB control block
D037 FFFE 0210 ( 8448) 386MAX UMB
D248 FFFA 0004 ( 64) 386MAX UMB control block
D24D FFFE 00C0 ( 5072) C:\404\4DOS.COM [2E ]
D30E FFFA 0004 ( 64) 386MAX UMB control block
D513 FFFE 0020 ( 512) 386MAX UMB
D554 D55C 0006 ( 96)
D53B D53C 04C0 ( 19664) Inv at 080A C:\UV\UV.COM [10 ]
D809 D53C 0002 ( 32)
D80C D814 0006 ( 96) [20 21 27 ]
D813 D814 01A0 ( 6864) Inv at 09C2 0:\UTILS\CTRLALT.COM [08 09 ]
D9C1 D814 0002 ( 32)
D9C4 FFFA 0004 ( 64) 386MAX UMB control block
D9C9 FFFE 0066 ( 1632) 386MAX UMB [27 23 24 ]
DA3D FFFA 0004 ( 64) 386MAX UMB control block
DA35 FFFE 0020 ( 512) 386MAX UMB
DA56 0000 05A8 ( 23168) free
DFF1 FFFD 1200 ( 73728) 386MAX locked-out area
F200 FFFF 0059 ( 1424) 386LOADED Driver
F25A 0000 05A4 ( 23104) free
F7F1 FFFD 0400 ( 16384) 386MAX locked-out area
FC00 FFFF 00AD ( 2768) 386LOADED Driver [4D 67 ]
FCAE FCBE 0006 ( 96)
FCB5 FCBE 0042 ( 1056) Inv at FCF9 C:\404\KSTACK.COM [16 ]
FCFB FCBE 0002 ( 32)
FCFB 0000 0004 ( 64) free
#D00 End of UMB Chain

```

```

0:\UDOS2\CHAP7> dev
NUL
CON
Block: 1 unit(s)
CON
Block: 1 unit(s)
CACHES$
386MAX$
EMMXXXX
CON
AUX
PRN
CLOCK$
Block: 3 unit(s)
COM1
.LPT1
.LPT2
.LPT3
COM2
COM3
COM4

```

In Figure 7-32 the output from UDMEM shows quite clearly that my device drivers really are resident in memory. Meanwhile, the output from DMA confirms that they are linked into the DOS device chain. For example, the first "CON" is ANSISYS, and the first block device is RAMDRIVE.SYS. Of course, the real test is that after loading RAMDRIVE.SYS and ANSISYS, I had an additional drive created by RAMDRIVE.SYS programs that assumed the presence of ANSISYS suddenly started producing reasonable output. And, of course, I had somewhat less memory.

It should be noted that some device drivers appear not to be properly loaded by DEVDIAG. These include some memory managers and drivers that use extended memory. For example, Microsoft's VMS driver, HIMEM.SYS, often crashes the system if you attempt to load it with DEVDIAG. This is an amazing coincidence, considering that HIMEM.SYS, if installed, is documented to require loading before any other memory manager.

Furthermore, while DEVDIAG\VDISK.SYS definitely works in that a virtual RAM disk is created, other programs that check for the presence of VDISK (such as protected-mode DOS extenders) often fail mysteriously when VDISK has been loaded in this unusual fashion. Again, VDISK is such an ill-behaved program that Microsoft lists it as being incompatible with Windows 3.1 under any circumstances, so discrepancies are to be expected.

Jay Lova has reported that a Triton SCSI driver appears to load through DEVDIAG without any detected errors, but that attempts to access the associated drive fail as if the driver were not there. Obviously, some mysteries remain to be solved here. However, in the vast majority of cases, DEVDIAG should give you no problems. The few cases we've run across that give trouble represent a limited collection of oddities, not representative of most device drivers.

For a better perspective on loading drivers, see the article by Tony Todd, "Installing MS-DOS Device Drivers from the Command Line," published in the British magazine *EM* (August 1989). For background on DOS device drivers in general, two excellent books are the classic *How to Use MS-DOS Device Drivers* (Second Edition, by Robert S. Taft, Reading, MA: Addison-Wesley, 1992) and *Writing DOS Device Drivers* (by Philip M. Adams and Charles E. Loeb, Englewood Cliffs, NJ: Prentice-Hall, 1990).

Many of the complexities of loading block devices—in particular, the importance of updating the CDS—become clear in the next chapter, where we discuss the DOS file system.



## The DOS File System and Network Redirector

by Jim Kyle, David Maxey, and Andrew Schulman

The file system is a vital and irreplaceable part of MS-DOS. While most successful PC software bypasses many of DOS's services and goes directly to the hardware to produce screen output or read the keyboard, few programs skip the DOS file system when it comes to reading and writing files. Even software like Microsoft's VLM 386 from Windows for Workgroups (WW 3.11) and "Chicago" DOS 7 Windows 4.0, which bypasses the file system code in MS-DOS, still closely simulate the behavior of DOS.

Actually, there are two DOS file systems. One, known as the FAT (File Allocation Table) file system from the name of its key data structure, is the logical structure that DOS uses for media such as floppy disks and hard drives; the FAT is probably the world's best-known DOS internal data structure, having entered popular culture through Peter Norton's book *Inside the IBM PC*. Even some books for non-programmers discuss the internal FAT structures as if use are needed for disk recovery.

The other file system, introduced in DOS 3.1, is known as the MS-DOS network redirector. Whereas most DOS programming interfaces consist of INT 21h or INT 2Eh functions that a program calls, the network redirector is instead a set of functions that MS-DOS itself calls. For example, when performing the open operation, DOS issues an INT 21h with AX=1116h. Any program can intercept INT 21h API calls and thereby make the contents of a network drive "real." When DOS tries to open a file, it calls `open`, calling on the program. The program can have the file opened (as if, for example, by sending a redirecting the request over a network to a file server). Another way to do this with NetWare servers prior to 4.0 allows to hook INT 21h directly and watch for a file-related call. However, as you'll see, there are some advantages to using the network redirector.

Thus, the network redirector is a set of hooks in MS-DOS that DOS uses for "trapping" a DOS directory function's call and routing from FAT systems such as network file servers and CD-ROM drives. Drives created with the network redirector do not require FATs or Disk Parameter Blocks. While networks are a particularly important part of the DOS file system, and one that discussion of DOS internals is quite relevant to, the network redirector is somewhat misnamed. It isn't just for networks; an internal network redirector is a mechanism, albeit a somewhat primitive one, for creating installable file systems.

All drives, whether FAT-based or non-FAT, have entries in a key DOS data structure called the Current Directory Structure (CDS). An important exception to this statement is Novell NetWare, which prior to version 4.0 bypasses the CDS. Many programs in this chapter manipulate the CDS in some way.

The CDS, together with many other DOS structures we discuss in this chapter, is shared with all programs. Since DOS is often thought of as a single-tasking operating system having one global structure, don't see this as a problem. However, as DOS is increasingly called upon to run multiple tasks, so-called "system" Windows, the assumption that only one program is using the CDS at a time becomes more and more unrealistic. Almost all DOS internal structures are global and non-rec-

rant. Finally, seriously restricting DOS's ability to perform true multitasking. That's one reason why multitasking versions of DOS, such as General Software's E-subbed DOS, avoid using such structures. See Mark Jones' "DOS Meets Real Time," *Unofficial System Programming*, February 1992. The way Windows employs instance data (Recall from Chapter 4 that Windows Enhanced Mode uses instance sets) to create the illusion of multiple DOSs (see Figure 1.9). Clearly, it would have been better if DOS supplied multiple DOSs (or an instance DOS) as the next best thing.

This chapter is concerned about DOS files, directories, and trees, and as in most such discussions, works with (even) with physical storage media and work out way to the directory structure seen by a typical DOS user. However, this chapter takes a somewhat different slant from most discussions of the DOS file system: even so, it emphasizes how DOS applies a logical ordering to physical media (a then it will be clear how to apply this same logical ordering to things other than hard drives and floppy disks). Any one of these sections (especially the ones on physical media) is more generic than the DOS notion of a drive is. It is not just for physical media or even RAM disks anymore.

There are several additional levels of complexity to the DOS file system. First, as with any modern operating system, some of the complexity is in the hardware. Disk caches such as SmartDrive introduce another level of complexity. There are also file systems. And, of course, the current FAT file system doesn't rest directly on top of the physical media, but instead goes through busk device drivers (see Chapter 7). So, there are "abstractions" to specify interfaces, such as the ability to read a given sector number, the stack-based "current" implementation of a file system as it sees it. And don't forget disk compression software such as Stacker and Multisoft's DoubleSpace. DOS 6.0. We discuss Stacker and DoubleSpace later in this chapter.

This chapter contains several hundred lines of sample programs giving a more of a cookbook approach to the file system. The chapter proceeds as follows: PHANTOM.C, a complete implementation of the DOS file system interface to treat a new drive. Readers of the first edition of this book will note that DOS 6.0 did not do this; we have completely rewritten the Phantom in C (instead of in assembly). And this is now a full-blown XMS RAM-disk, rather than a "proof of concept" toy. Other code in this chapter includes routines to:

- Detect DoubleSpace, Stacker drives, and RAM disks
- Get the compression ratio for a DoubleSpace drive
- Free up orphaned file handles
- Turn a cluster number into a file name, or vice versa
- Derive a filename from a file handle
- Increase the number of process file handles
- Determine the FILE\_S\* and BU\_FILE\_S\* values
- Set or turn off drive letters
- Walk the 12 bit and 16 bit File Allocation Table
- Walk the Current Directory Structure
- Walk the System File Table
- Get the true (canonical) name of a file
- List all the open files for any given process

What makes this a book on the computer system is more logical rather than the physical aspects of the DOS file system. But this book also includes a quick overview of the file system, followed by a book on the physical aspects.

## A Quick Overview of the System

To get a quick look at a piece of the DOS File System, from a perspective, we need to trace the significant actions that take place when DOS services a request to read from or write to a file. One good way to examine how services is to use the COPY command, which reads the content in one file and writes it to another.

One of our major tools for delving into the internal workings of DOS is INTRSPY, and it is program is what we use to see what happens inside the COPY command. Listing 8-1 shows the INTRSPY script we prepared. Basic and diverse tables of earlier versions of COMMAND.COM and the DOS service itself. This script includes such intermediary interfaces as INT 25h, 26h, INT 21h, 31h-12h, and the ROM BIOS INT 13h. With INTRSPY 2.0 we can even track right on the device driver, interrupt and strategy routines. See DISK.MR includes DD.MR from Chapter 5, Listing 5-2.

Incidentally, while we say that INT 13h is the ROM BIOS, disk interrupts, it is important to note that MS-DOS hooks INT 13h ahead of the BIOS, you can see this by running INTCHEAT 13/0/0/0/0 (see Chapter 6).

For a more sophisticated example, see Chapter 5, which uses INTRSPY to examine in detail the process of formatting a floppy disk.

### Listing 8-1: DISK.SCR

```

; DISK.SCR
; usage: cmdspy compile disk [drive] [command]
; examp.e: cmdspy compile disk c: command /c copy foo bar bar.foo > disk log

; DOS 4+/Compaq DOS 3.31+ >32M partition
structure big fields
sector (dword,hex)
num (word,hex)
addr (dword,ptr)

Intercept 21h
function 32h
on_entry output "2132: Get DPH drive " dx
on_exit output "2132: done"
function 3Ch
on_entry output "213C: Create File: " (ds dx->byte,asciiz,64)
on_exit output "213C: done, file is " ax
function 6Ch
on_entry output 216C Ext Open/Create (ds si->byte,asciiz,64)
output "AX=" ax " BX=" bx " CX=" cx " DX=" dx
on_exit output "216C: done, "
if (cflag==1) sameline "Error " cx
if (cflag==0) sameline "file is " ax
if (cx==1) sameline ", opened"
if (cx==2) sameline ", created"
if (cx==3) sameline ", replaced"
function 30h
on_entry output "2130: Open File: " (ds dx->byte,asciiz,64)
on_exit output "2130: done, file is " ax
function 3Eh
on_entry output "213E: Close File " bx
on_exit output "213E: done: File " bx
function 3Fh
on_entry output 213F Read File BX
on_exit output "213F: done: File " bx
function 40h
on_entry output 2140 Write File " bx (ds dx->byte,asciiz,cx)
on_exit output 2140 done: File " bx
function 44h
subfunction 00h
on_entry output 214400 IOCTL drive bl Attrs "
on_exit sameline dx
subfunction 09h
on_entry output 214409 IOCTL drive bl " Remote "
subfunction 0dh
on_entry output "21440d: IOCTL drive " bl
if (cl == 40h) sameline " [40: Set Device Parameters]"
if (cl == 41h) sameline " [41: Write Track]"

```

```

    if (cl == 42h) sameline " [42: Format and Verify Track]"
    if (cl == 60h) sameline " [60: Get Device Parameters]"
    if (cl == 61h) sameline " [61: Read Track]"
subfunction 0fh
    on_entry output "2144DF IOCTL Set Logical Drive " dl
    on_exit if (cflag == 0) sameline " ==> " al

```

## Intercept 25h

```

on_entry
    output "25: Abs Disk Read drv " al ", at sectr "
    if (cx == 0FFFFh)
        sameline (ds:bx->big.sector) ", " (ds:bx->big.num) " sectrs"
    if (cx != 0FFFFh) sameline dx ", " cx " sectrs"
on_exit if (cflag==1) sameline " [fail]"

```

## Intercept 26h

```

on_entry
    output "26: Abs Disk Write drv " al ", at sectr "
    if (cx == 0FFFFh)
        sameline (ds:bx->big.sector) ", " (ds:bx->big.num) " sectrs"
    if (cx != 0FFFFh) sameline dx ", " cx " sectrs"
on_exit if (cflag==1) sameline " [fail]"

```

## Intercept 13h

```

function 0 on_entry output "1300: Recalibrate drive " dl
function 1 on_exit output "1301: Disk system status " al
function 2
    on_entry
        output "1302: Read " al " sectrs: drv " dl ", head " dh
            , sectr " dl ", trk " ch " to " es " bx
    on_exit if (cflag==1) sameline " - FAILED (" ah ")"
function 3
    on_entry
        output "1303: Write " al " sectrs: drv " dl ", head " dh
            , sectr " dl ", trk " ch " from " es " bx
    on_exit if (cflag ==1) sameline " - FAILED (" ah ")"
function 4
    on_entry
        output "1304: Verify " al " sectrs: drv " dl ", head " dh
            , sectr " cl ", trk " ch
    on_exit if (cflag==1) sameline " - FAILED (" ah )"
function 5
    on_entry
        output "1305: Format " al " sectrs: drv " dl ", head " dh
            , sectr " cl ", trk " ch
    on_exit if (cflag==1) sameline " - FAILED (" ah )"
function 6 on_entry output "1306: Get drive params for " dl
function 0ch on_entry output "130C: Seek cy, ch drv dl " head " dh
function 0dh on_entry output "130D: Alternate reset drive " dl
function 10h on_entry output "1310: Test drive dl
function 15h on_entry output "1315: Get type drv dl
function 16h on_entry output "1316: Get media change drv dl
function 17h on_entry output "1317: Set type drv dl " al
function 18h on_entry output "1318: Set media type drv dl

```

```

include dd scr 21 22 3 24 5 76 27 28 9 , DD SCR does RUN, REPORT

```

This INTRSY script requires a command line with a drive letter and a DOS command. For example:

```

C:\UNDOC2\CHAP8>intrspy -R20480
E:\UNDOC2\CHAP8>cdpspy compile disk c command /c copy foo.bar bar.foo

```



In this example, FOO.BAR is a *text file* containing only the line "this is foo.bar". Figure 8-1 shows part of the INTRSPY output when copying FOO.BAR to BAR.FOO.

**Figure 8-1: INTRSPY Results for a DOS COPY**

```

216C: Ext Open/Create: FOO.BAR
AX=6C00 BX=0040 CX=0000 DX=0101
01 - media check
04 - input
1302: Read 09 sectors, drv 80, head 08, sector 08, trk 03 to 1785:0000
04 - input
1302: Read 03 sectors, drv 80, head 0A, sector 40, trk 36 to 1003:0000
1302: Read 01 sectors, drv 80, head 08, sector 40, trk 37 to 1063:0000
04 - input
04 - input
04 - input
216C: done, file is 0005, opened
214400 IOCTL drive 05 Attribs 0042
213E: Close File 0005
213E: done File 0005
01 - media check
01 - media check
216C: Ext Open/Create: FOO.BAR
AX=6C00 BX=0040 CX=0000 DX=0101
01 - media check
216C: done, file is 0005, opened
214400 IOCTL drive 05 Attribs 0042
213F: Read File 0005
04 - input
1302: Read 09 sectors, drv 80, head 01, sector 07, trk 04 to 1685:0000
1302: Read 01 sectors, drv 80, head 02, sector 42, trk 28 to 1003:0000
213F: done, File 0005
213E: Close File 0005
213E: done File 0005
216C: Ext Open/Create: BAR.FOO
AX=6C00 BX=0040 CX=0000 DX=0101
01 - media check
216C: done, file is 0005, opened
214400 IOCTL drive 05 Attribs 0042
213E: Close File 0005
213E: done File 0005
01 - media check
216C: Ext Open/Create: BAR.FOO
AX=6C00 BX=0021 CX=0000 DX=0112
01 - media check
04 - input
1302: Read 09 sectors, drv 80, head 02, sector 08, trk 04 to 1785:0000
08 - output
1303: Write 01 sectors, drv 80, head 06, sector 0A, trk 03 from 1223:0000
1303: Write 01 sectors, drv 80, head 02, sector 09, trk 04 from 101D:361C
08 - output
1303: Write 01 sectors, drv 80, head 08, sector 09, trk 05 from 101D:361C
08 - output
216C: done, file is 0005, replaced
214400 IOCTL drive 05 Attribs 0042
01 - media check
2140: Write File 0005this is foo.bar

2140: done: File 0005
214400 IOCTL drive 05 Attribs 0002
213E: Close File 0005
01 - media check
08 - output
08 - output

```

```

08 output
08 - output
03 - ioctl input
1303 write 01 sectors drv 80, head 08, sctr 06, trk 36 from 100C:0000
1303 write 03 sectors drv 80, head 0A, sctr 4B, trk 36 from 100B:0008
1303 write 07 sectors drv 80, head 08, sctr 48, trk 37 from 106B:0008
1303 write 05 sectors drv 80, head 02, sctr 09, trk 04 from 17F5:0000
1303 Write 05 sectors: drv 80, head 08, sctr 09, trk 05 from 17F5:0000
213E done: File 0005
2140 write File 0005 1
04 - input
1302 Read 01 sectors drv 80, head 0B, sctr 05, trk 46 to 1003:0000
2140 done: File 0001
2140 Write File 0001 file(s) copied
2140 done: File 0001

```

COPY uses the relatively new function 06h for opening and creating files instead of the older function 31h and 31h. When called to open BAR.FOO, this function calls block device driver function 1 Read, which in turn uses the BIOS INT 13h to read the root directory from the C drive. COPY then calls API function 4B to get the file's attributes and names the file.

Next, Figure 8-7 shows COPY opening the same file again. This time COPY reads the file's entire contents to a buffer before doing the file. Again, the command interpreter (COMMAND.COM) does not do the actual disk operation, just performing the open. The program then calls function 31h to do the actual disk operation. This calls BIOS to perform the actual work, though not directly, of course. Function 31h actually calls the block device driver Interrupt and Strategy routines, using device driver function 1 Read to the device driver, internally by the BIOS. For more information on how block device drivers fit into the DOS file system, see Robert S. Lee, *Writing MS-DOS Device Drivers*, second edition (Chapters 7 and 8).

When opening, reading from memory, the next step is to make function 06h by attempting to open the file. This is done by determining if the destination file BAR.FOO exists. If a call succeeds, indicating that the contents you want to see do exist in a directory, that information is made available to DOS buffers, so no guesswork/disk read is required. The success function 06h returns about the BAR.FOO is open, meaning that the file already exists from a previous test.

After success, API function 4C, a variable disk COPY uses the file then immediately reopens it for writing with the same IVT entry to re-open the file. This time you see several BIOS disk writes as the COPY function first erases some space the previous copy of BAR.FOO uses, updates both FATs, then modifies the contents of the file to a length of zero, which is getting with a new date-time stamp.

Next, COPY uses the file that was earlier read to write BAR.FOO, not as clearly seen in Figure 8-7. It first reads the data to be stored. Write function 0B, "2140 done" appears inmediately after entry to the program, with a completed 05 ends to a device driver of the BIOS. Another IOB 11 after with a check sum, indicates that writing has taken place. The attribute has change I from 00E2h to 00D2h. A check sum Write function 05 moved the data from COMMAND.COM's buffer area to the DOS buffers. As the data is not yet made out to the disk, 04h buffers play a crucial role in DOS file I/O, see BUFFERSC in Listing 8-8.

When function 31h finishes the file opening, DOS does all of the deferred writing. It allocates disk space, updates both FATs, writes the actual data with device driver function 8 Write, and finally updates the directory entry for BAR.FOO.

The final action shown is our trace file calls to the Write function to Handle 1 stdout to create the "File(s) copied" display.

At that point (copy) you file. And as noted earlier, even this inside view of the COPY command was quite superficial. We didn't get into what happens with disk compression, such as Stacker or DoubleSpace, or what happens if you're using a disk cache such as SmartDrive. Still, the INT13.PY

results provide a useful overview of the typical sequence of events involved in actual file I/O. Now let's look at the physical aspects of disk storage.

## The DOS File System

The starting point for the FAT file system is generally the physical disk and the drive mechanism itself. These marvels of mechanical precision convert a stream of information—represented as a sequence of bits—into a corresponding sequence of magnetic flux reversals that are processed on the surface of the disk.

Someone could write out the sequence of the methods involving this process, but a good deal of disk drive designers would read them. As programmers, we are more interested in how program-oriented descriptions of data are translated into the form of the actual disk hardware requires.

These translations occur in several steps. Programs organize data into a stream of bytes and store these streams in files which are later read back as streams. DOS translates file references to files into references to logical drive locations such as drive and cluster. The cluster is a part of the physical disk structure, but instead of merely a fiction introduced by the FAT code in DOS, DOS converts the cluster number into a logical sector number (LSN) for transmission to the disk device driver. If the device driver supports a physical disk, it translates the LSN into the more traditional oriented values of cylinder, head, and sector or transmission to the specific drive. The BIOS and its drive controller then translate those values into sequences of pulses that select the addressed drive, position the actuator to the desired cylinder or tracks, select the specified head, and begin reading from it when the correct sector is available. For a detailed examination of the BIOS and drive controller see Frank van Gorpwe's *The Unadorned PC*.<sup>1</sup>

## Surfaces, Tracks, and Sectors

One starting point to gaining an understanding of the DOS file system is the surface of the magnetic media on which is inscribed the sequential floppy diskette. The hard disk operates in much the same way, but with much greater precision.

In the earliest days of MS-DOS, the original IBM PC came equipped with a single-head, single-sided disk drive that had a storage capacity of 400K per disk. The head, which contacted the outer side of the diskette, was not raised or lowered, and its normal operating position, maintained by pressure of the head against the lower side of the diskette, was a 500-pressure gradient contact against the upper surface.

On the single-drive surface, the head contacted and later read back information from one of 40 concentric tracks. The head actuation mechanism was moved in or out to position the head accurately over the desired track. The track nearest the outer edge of the disk was designated as track 00, the nearest the hub hole, as track 39.

A vital feature of most of the single-head disk was a reference point to determine disk rotation. A sensor in the disk drive generated an index pulse each time this hole passed over it, and since the disk rotated at a constant speed of 300 RPM (200 milliseconds per revolution), the associated controller card could measure 100 sectors around the track in which to store data. These two sectors contained eight sectors per track, each sector with nominally 512 bytes of storage. Between sectors an address marker for sector identification codes helped the controller verify that it was well with the drive. Thus, each track contained 5\*8\*12 bytes of data, or 4,096 bytes, and the 40 tracks held a total of 163,840 bytes, or 160K.

Later, a single-drive system was supplanted by a two-headed, double-sided diskette that could read and write on both surfaces, immediately doubling the storage capacity to 320K per disk. MS-DOS 2.0 added a FAT sector to the format, bringing the storage capacity up to 360K. Later, at least 1.2MB drives—rotating at 300 RPM and reading 80 rather than 40 tracks—came along. The basic principles hold true, but there is, as well as our 3.5-inch units and today's huge hard drives,

In a few cases, the drive itself identifies storage locations in terms of which head is used, which track (or cylinder), and which sector of that track is read or written.

This has, however, made difficulty remembering a large collection of numeric values. Instead, we like to name things. It seems much simpler to remember that this text is stored in a file named `CHAPTER.DOC` than to try to remember that it is stored starting at sector 14, cylinder 93, head 5, of drive 3. That's part of what the DOS file system is all about. It seems easy to deal with our programs and data as named files. It may not be so easy to picture the job of translating these names into the sequence of numeric data that the machine requires. Since computers excel at dealing with numeric information, it is just another example of letting the computer do what it does best, so that humans can do what they do best.

Another aspect of the DOS file system extends this type of mapping to non-storage devices. RAM disks, for example, map a fraction of the stored file into fast volatile memory. DOS's simple I/O redirector facilities allow us to treat the screen and keyboard ports, CON, serial ports, COMs, and parallel ports, LPTs, as files. Drives connected with the DOS network or fraction can map a file system structure of sockets without the network to another machine, possibly running a completely different file system. The file system is other words not only simplifies access to hardware, but also provides unified access to otherwise disparate devices.

To deal with files with stored programs and data as named files, rather than forcing us to use physical head, track, sector addresses for every read or write action, the DOS file system maps these physical addresses into logical sector numbers and groups blocks of adjacent sectors into clusters for allocation to named files.

We examine these processes in detail shortly. But first, let's look at some special records that are not part of the file system, but without which the file system would not exist. These are the partition record, which exists in a specific physical drive as multiple logical drives, and the boot record, often called the boot sector, a sector of the DOS file record occupies a single sector, which controls the boot process each time you power up your system.

### Partition and Boot Records

The partition record came into use soon after hard disks became popular. The original purpose seems to have been to allow multiple operating systems, such as MS-DOS and UNIX, to exist on the same system, and to interface with each other. However, the capability quickly provided a way to deal with the 32-cylinder volume limit for disk drives that existed prior to MS-DOS 3.31, by allowing multiple 32-MPI logical drives on a single physical device. The presence of hard disk partitions shows that even a hard disk is just a logical construct, rather than a physical reality. Hard disk C may be just one subsection of the physical hard disk.

The boot record has been used with every file-based disk operating system. Its purpose is to control system specifications for that period of time when the full operating system has not yet been read into memory. The boot record sees to it that the operating system can be read from the disk.

Now, let's take a look at a partition record in the first sector of the first track under the first head of drive 80h (H:0:1:00:80). The I/O system supplies this record, sometimes also called the Master Boot Record (MBR). It establishes the physical limits on the logical drives and thus permits multiple logical drives to exist on a single physical drive. When you power up the system, the ROM BIOS reads the MBR into memory and transfers control to it. The code in this record, in turn, reads in the boot record for the currently specified bootable partition, then jumps to the code in that boot record.

Since the partition record is not inside any logical drive, and since DOS deals only with logical drives (the usual case is that each physical drive has only one logical drive, which occupies all available space), normally you cannot access this record. You can, however, read it with the BIOS disk read function. IN13h AH 02 as the following DEBUG script shows.

```

a
MOV AX,0201      ; AH=ReadSec, AL=number to read (1)
MOV BX,1000     ; buffer at ES:1000h, ES set by loader
MOV CX,0001     ; CH=cyl (0), CL=sector (1)
MOV DX,0080     ; DH=head (0), DL=drive (C:)
INT 13         ; call BIOS disk function
JMP 0100       ; provides place to set breakpoint

```

```

g 10e
d 1000:1200
q

```

The blank line after `JMP 0100` is essential: it signals `DEBUG` that the `A` assembly command is complete. To use this script for `DEBUG`, type it into a file `RPART.SCR` this type:

```
DEBUG < RPART.SCR > PART,CAP
```

This creates the same `PART.CAP` containing a hex dump of your `C:` drive's partition table. To read from your `D:` drive, you can change the value set for `DX` from `0080` to `0081`. Figure 8-2 shows an edited version of sample results from this script; we added spaces to create four-byte columns.

**Figure 8-2: A Partition Record**

```

000: FA 2B C0 BE  D0 8F C0 8E  08 08 00 7C  8B E0 F8 80      +.      |
010: F0 9F 00 7E  FC 89 00 01  F3 A5 E9 00  02 89 10 00
020: 8B 36 85 7E  F6 04 80 75  08 83 E1 10  12 F6 8B 37      6 " u " 7
030: 90 8F 8E 07  57 99 08 00  F3 A5 5E 8B  00 7C 8B 14      L M " "
040: 8B AC 02 80  05 00 88 01  02 CB 13 75  09 28 C0 CB      M " 3 U u
050: 13 40 74 19  EB 10 BE FE  7D AD 30 55  4A 75 14 BE      | 6 " 6 " 6
060: 8E 07 EA 00  7C 00 00 88  36 87 7E E8  0A 8B 36 89      6 " " " |
070: 7E E8 04 88  36 88 7E AC  0A C0 74 FE  BB 07 00 84      6 " " " |
080: 06 C0 10 88  72 EE 7F 89  7E A7 7E C8  7E 00 0A 89      nval id Partition
090: 8E 76 61 8C  69 44 20 50  61 72 74 69  74 69 6F 6E      Table Error L
0A0: 20 54 61 62  6C 65 00 09  0A 65 72 72  6F 72 20 4C      oading Operating
0B0: 6F 61 64 69  6E 67 20 4F  70 65 72 61  74 69 6E 67      System Missing
0C0: 20 53 79 73  74 65 80 00  00 0A 40 69  73 73 69 8F      g Operating Syst
0D0: 67 20 4F 70  65 72 61 74  69 6E 67 20  53 79 73 74      em
0E0: 65 60 00 00  00 00 00 00  00 00 00 00  00 00 00 00
0F0: 00 00 00 00  00 00 00 00  00 00 00 00  AA 55 00 00
100: 00 00 00 00  00 00 00 00  00 00 00 00  00 00 00 00
      lines omitted, all zeroes
1B0: 00 00 00 00  00 00 00 00  00 00 00 00  00 00 80 01
1C0: 01 00 06 08  B1 FE 11 00  00 00 56 63  02 00 00 00      yC
1D0: 00 00 00 00  00 00 00 00  00 00 00 00  00 00 00 00
1E0: 00 00 00 00  00 00 00 00  00 00 00 00  00 00 00 00
1F0: 00 00 00 00  00 00 00 00  00 00 00 00  00 00 5A AA      . .

```

The hex 254 bytes (4 to index) give the signature `AA55` in offset 0Eh for the record's table entry (hex 0D00 to 0E00), which of the 32 partitions is used for the startup process. The remaining 255 bytes (up to offset 0Fh) are a table of system ID, 16-byte entries (each of 5 bytes) as the bits of a special partition. In Figure 8-2, one such entry starts with the byte 80h at offset 1Bh.

Microsoft's `FDISK` uses the layout of each entry (see the `PARTENTRY` structure in the `MS-DOS 3.86 source`, `BIOS\src`). This is the arrangement for each of the 16 entries:

```

struct PartEntry {
    char BootableFlag,      /* 80h = bootable partition, else 00h */
    char StartHead;        /* # starting head number */
    char StartSector,      /* # Bits 0-5 are start sector, 6-7 cy */
    char StartCyl,        /* # Low 8 bits of start cyl, 8-9 to prev */
    char SystemID,        /* # Encodes file system type, see below */
    char EndHead;         /* # ending head number */

```

```

char EndSector,      /* ending sector and hi bits of cyl. */
char EndCyl,        /* low bits of ending cyl number */
unsigned long AbsBegin, /* nbr, relative to disk, of first sec */
unsigned long SectCount, /* total number of sectors in partition */
};

```

The `System` field byte can hold any value, since non-MS-DOS operating systems also use this partition table. As of version 5.0, MS-DOS recognizes the following values:

- 00h = Unknown type or unused entry
- 01h = MS-DOS, 12-bit FAT
- 04h = MS-DOS, 16-bit FAT, partition < 32 MB
- 05h = MS-DOS, extended partition
- 06h = MS-DOS, 16-bit FAT, partition >= 32 MB

A good reference provides values of other operating systems in Chapter 16 of his *DOS Internals*.

The cylinder shows there is no corresponding partition 1 (Figure 8-2) only one of the sixteen possible cylinders is non-zero, because this drive was partitioned as a single 70-megabyte logical volume and not as a set of 11H cylinders that this drive is capable of. The value of 11Hh says that the partition occupies all cylinders at offset 1000 specifies sector 1, and that at 13Hh, together with the high two bits of 100h, indicates a track 00, the 00h at offset 1C20h shows that this is a DOS volume exactly 16 of 16 megabits larger than 32 megabytes. The next three bytes specify the ending head sector and the partition is 16 sectors (00000016h) starting at offset 1C00h means that the partition begins with a 729 sector from the start of the physical disk, and the final value 0001263560h at offset 1C40h is the total sector count for the volume.

Partitions are created with the FDISK utility. Commented C source code for an FDISK utility, along with code for FORMAT, CHKDSK, SYS, DISKCOPY, and other programs, is included with the excellent Unix-NTDK available from General Software, Redmond, WA.

### The Boot Record and BIOS Parameter Block (BPB)

The boot record, which Microsoft calls the *bootstrap record* in the MS-DOS 5.0 documentation, occupies the first sector of a DOS bootable partition. FDISK will not create more than one DOS bootable partition on any physical drive. The disk or the partition record reads the boot record into memory for the system to boot record starts DOS disk volume with the DEBUG command to load the first sector into memory. For instance, `C:\>02001` at the DEBUG prompt reads one sector of start sector 1, first sector 0, and disk C: 2 to offset 1000 (1000) DEBUG, and the rest of DOS start cylinder sector with zero elsewhere, the first sector is sector 1 (Figure 8-3 shows how to view DOS boot records with DEBUG).

**Figure 8-3. Using DEBUG to View DOS Boot Records**

```

C:\>UNDOC2\CHAP8>debug
-L 100 2 0 1
-d 100 200
81E2:0100 EB 00 90 53 54 41 43 40-45 52 20 00 02 10 01 00    ...STACKER .....
81E2:0110 02 00 02 00 00 F8 66 00 3F 00 0F 00 00 00 00 00    .. 2 .....
81E2:0120 4B 25 03 00 00 14 00 00-00 01 09 00 20 47 32 00    M.....62..
81E2:0130 00 0E 00 ED 00 CB 00 00-00 88 00 00 00 00 00 00    .....
81E2:0140 00 53 54 41 43 56 4F 4C-20 30 30 43 79 00 AB      STACVOL 000C3..
..
.. etc
.. that's a Stacker drive; let's look at its host, a normal DOS drive
-L 100 5 0 1
-d 100 200
81E2:0100 EB 3C 90 4B 53 44 4F 53 35 2E 30 00 02 04 01 00    <.MSDOS5.0.....
81E2:0110 02 00 02 00 00 F8 7A 00-11 00 0C 00 11 00 00 00    .. 2 .....
81E2:0120 9F E7 03 00 80 00 29 E7-11 51 1B 4E 4F 20 4E 41    .....3...Q.NO NA

```

```

81E2:0130 4D 45 20 20 20 20 46 47 54 51 56 20 20 20 FA 33 ME FAT16 3
81E2:0140 C0 88 00 BC 00 7C 76 07 BB 78 00 56 C5 37 1E 56 . .] x 6 7.0
; --
;u 0100 102
81E2:0100 E93C JMP 015E ; bsJump
81E2:0102 90 NOP
; --
0103h db "MSDOS5 0" ; bsDevName
;;; beginning of BPB
0108h dw 0200h ; bsBytesPerSec
010Ch db 04h ; bsSecPerClust
010Eh dw 0007h ; bsResSectors
0110h db 02h ; bsFATS
0111h dw 0200h ; bsRootDirEnts
0113h dw 0000h ; bsSectors (0, see bsHugeSectors)
0115h db FBh ; bsMedia (FBh = hard disk)
0116h dw 007Ah ; bsFATsecs
0118h dw 0011h ; bsSecPerTrack
011Ah dw 000Ch ; bsHeads
011Ch dd 00000011h ; bsHiddenSecs
0120h dd 0001E79Fh ; bsHugeSectors
;;; end of BPB
0124h db 80h ; bsDriveNumber (80h = first hard disk)
0125h db 00h ; bsReserved1 (used during boot)
0126h db 29h ; bsBootSignature (29h = extended BPB)
;;; start of extended BPB (see Media ID)
0127h dd 1B5111E7h ; bsVolumeID (midSerialNum)
012Bh db "NO NAME" ; bsVolumeLabel (midVolLabel)
0136h db "FAT16" ; bsFileSysType (midFileSysType)
;;; end of extended BPB

```

The first three bytes of Figure 8-3 (MP) are the actual code you can unassemble, of course, that first verifies that the disk is indeed a system disk. It is the code reads in the BIOS file, which then takes over the startup process.

Immediately after the JMP instruction is a set of data items that Microsoft's master documenter in the DOS programmer's reference (see BIOS, SIC 10R) has set included a structure called the BIOS Parameter Block (BPB), documented in the device driver chapter of the DOS programmer's reference.

As Figure 8-3 shows, the BPB contains information about the storage medium, such as the bytes per sector, sectors per cluster, number of FATs and root directory entries, and so on. As noted earlier, however, DOS may overwrite some of the values in a BPB so the BIOS that DOS runs in its memory doesn't necessary watch a BPB on a disk. For example, the number of FATs in floppy and hard disks is 2, no matter what the disks use BPB entries. For further discussion, see Chapter 15, *DOS Internals*, Chapter 16.

Note the presence of what is called an extended BPB. You can retrieve this data with the generic IOCTL Get Media ID call (see 21h AX, 4400h, CX, 0800h), which is equivalent to a slightly preferred IOCTL disk ioctl (21h AX, 0900h). The inclusion of a second, and by the way, extended BPB means that a volume can be managed by disk with a special attribute, i.e., root directory. In MS-DOS 4.0 and higher, you can set an Equivalency Label with the IOCTL Set Media ID call. Note, however, that this call only works for media that in fact does have an extended BPB. For disks formatted under MS-DOS 3.0, for example, Get Media ID just is a raw code to access the data, and they must retrieve it by the volume label using an old method if it requires an IOCTL.

Don't be surprised further extended the BPB to a structure called the MDPBPB, the MDP stands for MagnaDisk, see the DoubleSpace discussion later in this chapter.

### Logical Sector Numbers and the Cluster Concept

The first step toward simplifying the head-track-sector number sequence was to recognize that there was a more efficient way of uniquely specifying every sector on a disk unit with a single number, rather than with three.

DOS assigns numeric numbers to the sectors in logical sequence. That is, the first sector (S) of the first logical track (T) under the first logical head (H) — which in the fully hardware-oriented scheme, takes into account the position of the offset or head 1 word of the H+T+00 S=1 — becomes Logical Sector Number (LSN) 001. If we let every of the PC start numbering physical sectors from 1, not 0. The calculations used to determine the LSN work with values that are relative to the start of the logical drive from the first surface of the physical drive. The rest of the way around that first track on the same surface, the numbers follow in sequence.

Let's reverse the LSN map to the other surface of the disk. For a 360K diskette, with no partition record and nine sectors per track on both sides, LSN 10 would be H=1, T=00, S=1. After all sectors on track 0 of the second side are accounted for, the numbering returns H=0, T=01, S=1 for track 1 of the first side, which becomes LSN 19.

It's more difficult to calculate the exact translation for a floppy, but the essential point is that you can convert any head-track-sector reference into a unique LSN if you know how many sectors are on each track, how many tracks the disk has, and where the logical volume begins with respect to the physical one. You can reverse the process by reverse translation.

With extended DOS partitions, both translations can be difficult because no clean way exists with extended DOS to determine the start location of the partition, other than by reading the partition record into the DOS file system and finding the data found in the appropriate entry there. The problem with extended sectors is determining which file applies to the extended partition with which you are dealing. For more information on extended DOS partitions, see Chapter 16 of *Good Stuff Happily DOS Edition*.

File-based storage units, which may contain hundreds of thousands of 512-byte sectors, the LSN is too cumbersome. The DOS typically needs to allocate disk space and to access files. This emerged the idea of clusters. DOS inherited this idea from the older CP/M operating system, although CP/M see the term as a word. A cluster is simply a group of adjacent sectors that are always assigned to a file. If the file is one byte, it gets a whole cluster anyway. This solves several problems and creates one new one.

Most files are multiple sectors. If storage is by address, DOS overhead in allocating and tracing disk space over DOS performs these actions a fraction less frequently than it would if it allocated space one byte at a time. Multiple sector clusters also serve to speed up disk access by reducing (though by no means eliminating) the extent to which the file can be one scattered all over the drive. Even if it is two clusters in the same adjacent track, other than a cost in time each cluster all the sectors are together. Since sectors are the unit of disk I/O delay, this improves overall system performance.

One major disadvantage of clusters is that when disks are more than one sector per cluster they are less than 100 percent disk space occupied by files. If the disk space file size is always a multiple of 8 sectors per file, for example, but each 512 bytes per sector and 8 sectors per cluster, then the inefficiency of the disk space is one sector per file whose size is a directory listing is one byte. Not exactly a peanuts issue. This wasted space is known as cluster overhead.

Not having a cluster of a few sectors. Some RAM disks such as Microsoft's RAMDRIVE.SYS use one sector clusters. Lower capacity diskette formats use a cluster of only two sectors. 1.2 MB and higher capacity diskettes use one sector clusters. Hard disks for the most part use either 4 sector or 8 sector clusters. Some optical drives use larger clusters as an alternative to larger sector sizes, to deal with big wastes of space under older versions of DOS. DOS has a built-in limit of 128 sectors per cluster in MS-DOS 5.0, some PCB handling breaks down under this condition, so in practice 64 sectors per cluster is tops.



One way that DoubleSpace compresses disks is by allocating a *variable* number of sectors per cluster. Rather than converting from clusters to sectors via a fixed formula, DoubleSpace uses a lookup table called the MIB (Master Index Block). We discuss DoubleSpace in greater detail later in this chapter.

To tell which clusters are in use, and which are available for assignment, DOS uses the famous File Allocation Table (FAT)—with one entry per cluster. The cluster structure is the backbone of the FAT file system. If a subdirectory of data on the affected disk can't be found,

### The File Allocation Table (FAT)

The FAT is always located near the front of each disk volume, generally immediately after the boot record. The FAT is a sequence of what would normally be successive boundaries of the first sector after the boot record. DOS normally contains two copies of the FAT in case of hard-disk errors—of course, real errors. DOS may successfully write both copies, or some space is allocated or released, but only needs to read one of them. DOS ignores an error reading the first copy of FAT, successfully read the second copy. Microsoft uses three FATs on its brand-name Disk BASIC (1979); this is where the multiple FAT idea, and in fact FAT itself, comes from. For some of the goals of FAT, see Tim Paterson, "A Disk Look at MS-DOS," *Byte*, June 1983.

The FAT is a series of cluster numbers. That is, you use a cluster number as an index into the FAT. The value at FAT[cluster number] is either another cluster number or an end-of-the-indicator. These cluster numbers indicate the location of disk files, directories, and free space.

Sometimes each element of the FAT may be a 16-bit cluster number, but it may not match for smaller the hardware FAT size may not conveniently sized 12-bit numbers.

For DOS 3.x and higher, the top 4 bits of the highest cluster word of each drive parameter block tell which FAT size is used for a specific volume. DOS 4.0 added a field for an extended BPB, containing an 8-character file-system identifier, which can be "FAT12" or "FAT16" with the unused three bytes padded with space characters (see *Windows* type Figure 8-3). You can access this field with FN 21h AX=4401h CX=0806h Get Media ID).

Because of the FAT16 identifier, it's technically possible to have a 16-bit FAT even for disks with 11-bit cluster sizes. However, the DOS kernel uses pointers to FAT entries when these are sufficient to identify a media's highest cluster number. Although the Sector ID field in the parameter record identifies which parameter uses a 12-bit or 16-bit FAT, DOS doesn't use this information. The reason for this seemingly stupid sight is that floppy and other removable media are formatted with pointers that still require FAT to be usable under DOS. So in practice, it's difficult to just check the top four bits of the highest cluster (bits 12-15) to determine the FAT size.

Even with the huge volume sizes that DOS 4.x can support, the FAT entry size of DOS itself never exceeds 48 bits, despite occasional claims to the contrary. What does increase as the volume size grows is the cluster size, with the mathematics noted earlier. The maximum FATs Hooks are available to permit creation of custom file systems using other FAT sizes. A third party vendor perhaps could write a FAT 24 or FAT 32, but DOS itself recognizes only the 12-bit and 16-bit sizes.

Each element of the FAT, whether 7 or 16 bits long, corresponds to a single cluster of the drive's storage space. The first two locations, which would refer to cluster 0 and cluster 1, instead hold media information. The first byte of the FAT indicates the Media Code, if the drive serves removable media, and a variable correlation with the drive's own type. The remaining 16 or 24 bits of the first two entries are normally set to zeroes and remain unused. It appears that DOS may use cluster 1 as a temporary marker while building an allocation chain for a file, and it uses cluster 0 as a shortcut reference to the root directory for any drive where defining the directory entry for a subdirectory from the root

Cluster 2 is the first one usable for data. Since both copies of the FAT and the volume's root directory area precede this space, DOS must calculate the LSN for cluster 2 from the values provided in the DPB. Note, though, that DOS is not consistent in performing this calculation, as a result of a few directory files not containing sector values that are exact multiples of 16 entries; things become confused. DOS rounds down when doing the calculation, while MSDOS.SYS rounds up.

Given the LSN for cluster 2, you know the LSN at which any cluster starts; then multiplying that cluster number by 16 (the sectors per cluster), then adding the known LSN for cluster 2, that's your LSN. DOS file status cluster numbers take it from directory entries into the LSNs that disk devices need to receive. As noted earlier, the LSN for cluster 2 is calculated from values in the DPB. But we're getting ahead of ourselves; DPBs are explained shortly.

The also-mentioned `fat` command tells whether the corresponding cluster is in use or not, and how. It's a cluster status essential information about the file that is using it. A zero indicates that the cluster is free and can be reused. A value of 1 should never occur, although you can trap such a value on disk 1 from the root of C: as a result of an `INT 21h` function or as single stepping through an `INT 21h` function and have a low enough setting for `BUFFERS`.

The status codes possible range from 118h-111h for 12-bit FATs, or 118h-111h for 16-bit FATs. (111h, that is, the cluster is the last one in the file. 117h marks a bad cluster; 110h through 116h are reserved, meaning that they are not used and quite possibly never will be. Any other value indicates that the corresponding cluster is continued in the cluster having that value, next cluster = 1 + [cluster].)

The cluster 1 sector 1110h is highest cluster number, because 2 is the first valid cluster number, and because each file (no matter how small) occupies at least one cluster. There can be at most 1110h - 1111h = 65535 clusters and one sector per DOS volume. 65535 is not a large number.

List 8-2 shows some sample printouts from `FAT C`, which prints out the FAT chain for any drive and cluster. Lists 8-3 and 8-4 are also on the same line. For example, the following shows that the file starting at cluster 1100 on drive C occupies 58 clusters, distributed in nine different places on the disk.

```
C:\WINDOWS\CHAP8>fat c, 7470
7470-7486 (17)
7491-7492 (2)
7494-7502 (9)
7512-7516 (5)
7544-7545 (2)
7554-7556 (3)
7683-7715 (33)
7729-7746 (18)
7752
88 clusters in 9 groups
(10% fragmentation)
```

But you can examine a file's sectors to cluster 7470 on drive C. You'll see how filenames are mapped to clusters in the `FILE INFORMATION` section of `NTFSUTIL.C` in Listing 8-5.

## Listing 8-2: FAT C

```
/*
FAT C -- Given drive and cluster number, print FAT chain
Andrew Schulman, July 1993
*/

#include <stdlib.h>
#include <stdio.h>
#include <ctype.h>
#include <dos.h>
#include "diskstuf.h"

void fail(const char *s) { puts(s); exit(1), 3
```

```

main(int argc, char *argv[])
{
    DPB far *dpb;
    WORD prev_cluster, cluster, num_clusters = 0, num_groups = 0,
    WORD start_cluster,
    int drive,

    if (argc < 3)
        fail("usage: fat [drive] [cluster]");
    drive = toupper(argv[1][0]) - 'A';
    cluster = atoi(argv[2]); // decimal, not hex

    if (! (dpb = get_dpb(drive + 1))) // see DISKSTUF.C (listing 8-4)
        fail("can't get DPB");
    start_cluster = cluster,
    // the following works because get_fat_entry() returns
    // 12-bit EOF in 16-bit form
    #define EOF(cluster) ((cluster) == 0xffff) // end of file
    while (! EOF(cluster))
    {
        num_clusters++;
        prev_cluster = cluster;
        cluster = get_fat_entry(drive, dpb, cluster),
        #define CONTIGUOUS(x,y) ((x) == ((y) + 1))
        if (! CONTIGUOUS(cluster, prev_cluster))
        {
            int num_clust = prev_cluster + 1 - start_cluster,
            num_groups++;
            if (num_clust > 1)
                printf("u-lu (%u)\n",
                    start_cluster, prev_cluster, num_clust),
            else // only one cluster
                printf("%u\n", start_cluster);
            start_cluster = cluster,
        }
    }
    printf("\nu clusters in %u groups\n", num_clusters, num_groups),
    if (num_clusters > 2 && num_groups > 1)
        printf(" (%u% fragmentation)\n", (num_groups * 100) / num_clusters),
    return 0;
}

```

FATC invokes the header file DISKSTUF.H (listing 8-3) and uses the `get_dpb` and `get_fat_entry` functions from DISKSTUF.C (listing 8-4). We use DISKSTUF again later in this chapter. All the hard work in FATC is doing `get_fat_entry`, although hard work is a loose statement when working with 16-bit FATs; all the difficulty is in handling 12-bit FATs. FATC depends on `get_fat_entry` to return 12-bit EOF markers in 16-bit form. Because `get_fat_entry` masks the difference between 12-bit and 16-bit FATs, FATC doesn't know or care what type of FAT it's dealing with.

### Listing 8-3. DISKSTUF.H

```

/*
DISKSTUF.H -- Some functions and structures for low-level disk access
Andrew Schulman, July 1993
*/
#ifndef DISKSTUF_H
#define DISKSTUF_H

typedef unsigned char BYTE;
typedef unsigned short WORD;
typedef unsigned long DWORD;

```

```

typedef char *STRING,
extern void fail(const char *s); // application must define
#pragma pack(1)
typedef struct {
    BYTE name[8], ext[3], attrib,
#ifdef OS2
    BYTE reserved[8],
    WORD ea_handle, // OS/2 handle to extended attributes (EAs)
#else
    BYTE reserved[10];
#endif
#ifdef FAT
    WORD time, date, cluster;
    DWORD size;
} DIR_ENTRY,
#define VOLUME_ATTR 0x0B
#define DIRECTORY_ATTR 0x10
#ifdef MK_FP
#define MK_FP(seg,ofs) \
    ((void far *)(((DWORD)(seg) << 16) | (ofs)))
#endif
typedef struct dpb { // Disk Parameter Block
    BYTE drive, unit;
    WORD bytes_per_sect;
    BYTE sectors_per_cluster; // plus 1
    BYTE shift; // for sectors per cluster
    WORD boot_sectors,
    BYTE copies_fat,
    WORD max_root_dir, first_data_sector, highest_cluster,
    union {
        struct {
            BYTE sectors_per_fat,
            WORD first_dir_sector, // root dir
            void far *device_driver,
            BYTE media_descriptor, access_flag;
            struct dpb far *next;
            DWORD reserved,
            } dos3,
        struct {
            WORD sectors_per_fat; // WORD, not BYTE!
            WORD first_dir_sector,
            void far *device_driver,
            BYTE media_descriptor, access_flag,
            struct dpb far *next;
            DWORD reserved,
            } dos4,
        } vers,
    } DPB,
#pragma pack()

DPB far *get_dpb(int drive),
int _dos_driverremovable(int drive);
int _dos_getdrivemap(int drive),
char far *truenam(char far *s, char far *d);
WORD get_fat_entry(int drive, DPB far *dpb, WORD cluster);
int read_sectors(int drive, BYTE far *buf, int sectors, DWORD first_sector);
#endif /* DISKSTUF_H */

```

[ ] DISKSTUF.H Listing 8-4 get\_dpb checks for removable media and installs an INT 24h critical error handler. This is because the DOS C++ DPB function INT 21h AH=32h hits the disk. In the absence of an INT 24h critical error handler like the one provided in Listing 8-4, calling Get DPB for drive A: for example, can produce an annoying "Not ready reading drive A / Abort, Retry, Fail?"

message. Furthermore, asking Cact DFB for drive B, if the drive is currently assigned to drive A produces the equally annoying "Insert diskette for drive B, and press any key when ready" message. This issue is discussed a length in, as usual, Chapter 16 of *Ceoff's Cappel's DOS Internals*.

#### Listing 8-4: DISKSTUF.C

```

/*
DISKSTUF.C -- Some functions and structures for low-level disk access
Andrew Schulman, July 1993
*/

#include <stdlib.h>
#include <stdio.h>
#include <ctype.h>
#include <dos.h>
#include "diskstuf.h"

int _dos_driverremoveable(int drive)
{
    int retval;
    __asm mov ax, 440Bh
    __asm mov bl, byte ptr drive
    __asm int 21h
    __asm jnc ok
    return 0; // treat error as non-removeable
ok:
    __asm mov retval, ax
    return (!retval); // turn around so non-removeable = 0
}

int _dos_getdrivemap(int drive)
{
    __asm mov ax, 440Eh
    __asm mov bl, byte ptr drive
    __asm int 21h
    __asm jc error
    __asm cmp ax, 0 // only one drive number
    __asm je error
    __asm xor ah, ah
    __asm mov word ptr drive, ax // active drive number
error:
    return drive;
}

typedef struct {
#ifdef __JRB0C__
    unsigned short bp,d1,s1,ds,es,dx,cs,bx,ax,
#else
    unsigned short es,ds,d1,s1,bp,sp,bx,dx,cs,ax, /* PUSHA order */
#endif
    unsigned short ip,cs,flags;
} REG_PARAMS;

/*****
typedef void interrupt (far *INTVECT)(),
static volatile int failed = 0;
static INTVECT old_int24 = (INTVECT) 0;
void interrupt far crit_err(REG_PARAMS regs) // INT 24h handler
{
    REG_PARAMS *pregs = &regs;
    pregs->ax = 3;
    failed++;
}
void watch_crit_err(void)
{
    failed = 0;

```

```

old_int24 = (INTVECT) _dos_getvect(0x24);
_dos_setvect(0x24, (INTVECT) crit_err);
)

void unwatch(crit_err_t void) { _dos_setvect(0x24, old_int24); }
int get_failed(void) { return failed; }
void reset_failed(void) { failed = 0; }
/*****
DPB far *get_dpb(int drive)
(
    DPB far *dpb = (DPB far *) 0;
    // if drive removable and not mapped, fail
    if (!_dos_driverremovable(drive))
        if (_dos_getdrivemap(drive) != drive)
            return (DPB far *) 0;
    // install temp crit-error handler for duration of 21/32 call
    watch_crit_err();
    // call newly-documented DOS Get DPB function
    _asm push ds
    _asm mov ah, 32h
    _asm mov dl, byte ptr drive
    _asm int 21h
    _asm mov dx, ds
    _asm pop ds
    _asm cmp al, 0ffh
    _asm je fin1
    _asm mov word ptr dpb+2, dx
    _asm mov word ptr dpb, bx
fin1
    // remove temporary crit-err handler
    unwatch_crit_err();
    // "failed" is set inside Get DPB by crit-err handler
    return failed? ((DPB far *) 0) : dpb;
)

char far *truename(char far *s, char far *d) // get canonical pathname
(
    /* INT 21h AH=0Dh doesn't like leading or trailing blanks */
    char far *s2;
    while (isspace(*s)) s++; // ltrim
    s2 = s; while (*s2) s2++; s2--; // go to end
    while (isspace(*s2) *s2-- = 0; // rtrim
    /* Apparently some versions of DR DOS insist on ES:DI != DS:SI */
    _asm push di
    _asm push si
    _asm push ds
    _asm push es
    _asm les di, d
    _asm ds si, s
    _asm mov ah, 0Dh
    _asm ni 21h
    _asm pop es
    _asm pop ds
    _asm pop si
    _asm pop di
    _asm jc error
    return d;
error:
    return (char far *) 0;
)

typedef struct {
    DWORD diStartSector;
    WORD diSectors;
    BYTE far *diBuffer;
}

```

```

) DISKIO,
int read_sectors(int drive, BYTE far *buf, int sectors, DWORD first_sector)
{
    DISKIO diskio, far *pdiskio = &diskio,
    diskio.dStartSector = first_sector,
    diskio.dSectors = sectors;
    diskio.dBuffer = buf;

    // INT 25h/26h are obsolete, should instead use generic
    // IOCTL functions 21/440D/0861 and 21/440B/0841.
    __asm push ds
    __asm mov al, byte ptr drive
    __asm lds bx, pdiskio
    __asm mov cx, 0FFFFh
    __asm mov dx, 0
    __asm int 25h
    __asm jc do_fail
    __asm popf // INT 25h, 26h leave flags on stack
    __asm pop ds
    return 1,
do_fail
    __asm popf
    __asm pop ds
    return 0,
)
// get_fat_entry() hides differences between 12-bit and 16-bit FATs
// returns 12-bit FAT EOF indicators (e.g., 0xFFD) in 16-bit form (0xFFFFD)
WORD get_fat_entry(int drive, DPB far *dpb, WORD cluster)
{
    static WORD *fat_sect = (WORD *) -1,
    static WORD prev_drive = (WORD) 1,
    static DPB far *prev_dpb = (DPB far *) -1L,
    static WORD prev_sect = (WORD) -1,
    WORD first_fat_sect, entry_per_sect, sect, fat_size,
    // One-time initialization of buffer, allocate 2 sectors because
    // this may be needed for 12-bit FAT if a FAT entry slops over
    // two sectors. [Yuch3]
    if (fat_sect == (WORD *) -1)
        if (! (fat_sect = (WORD *) malloc(dpb->bytes_per_sect * 2)))
            fail("insufficient memory");
    if (_osmajor >= 4)
        first_fat_sect = dpb->vers.dos4.first_dir_sector -
            (dpb->copies_fat * dpb->vers.dos4.sectors_per_fat),
    else
        first_fat_sect = dpb->vers.dos3.first_dir_sector -
            (dpb->copies_fat * dpb->vers.dos3.sectors_per_fat),
    fat_size = (dpb->highest_cluster >> 12 == 0) ? 12 : 16,
    if (fat_size == 12)
        sect = (1 - first_fat_sect + (((cluster + 3) / 2) / dpb->bytes_per_sect));
    else
    {
        entry_per_sect = dpb->bytes_per_sect / 2;
        sect = first_fat_sect + (cluster / entry_per_sect);
    }
    // Don't reread if same sector as last time (assumes of course
    // that some TSR hasn't modified the sector in the meantime)
    if (! (drive == prev_drive && dpb == prev_dpb && sect == prev_sect))
    {
        int num_sect = (fat_size == 12) ? 2 : 1; // for possible s.o.p
        if (! read_sectors(drive, (BYTE far *) fat_sect, num_sect, sect))
            fail("can't read FAT"); // could try FAT #2
    }
    prev_sect = sect, prev_drive = drive, prev_dpb = dpb,

```

```

if (fat_size == 12)
{
    BYTE *fat = (BYTE *) fat_sect;
    WORD ofs = (cluster * 3) / 2) * dpb->bytes_per_sect;
    WORD retval = *((WORD *) &fat[ofs]);
    if (cluster & 1) // odd cluster #
        retval >>= 4; // take top 12 bits
    else // even cluster #
        retval &= 0x0FFF; // take bottom 12 bits
    // return 12-bit FFD-FFF as 16-bit FFFD-FFFF
    return (retval > 0x0FFF) ? (retval < 0x0000) ? retval,
}
else // gosh, 16-bit FATs are easy
    return fat_sect[cluster * entry_per_sect];
}

```

Inside `fat.c`, `get_fat_entry()` and `get_fat_entry()` (DISKLIB.FA also contains the function `read_sector`, which uses INT 13h's new-style first introduced in Compaq DOS 3.31) INT 25h and 26h calls `FDISK.Absorb` Disk Read and Write functions. According to the MS-DOS programmer's reference, the generic IOCL functions INT 21h AX=44010h CX=0801h Read Track on Logical Disk and CX=0801h Write Track on Logical Disk are now the preferred functions for accessing disks via the BIOS. For example, FORMAT relies heavily on generic IOCL. The disk read function IOCL and also largely replacement need for directly messing with BIOS INT 13h calls into INT 25h/26h via the generic IOCL disk functions just call down to the block device driver, which in turn does whatever is needed to read the disk. As shown by running the FAT program under IN-GENY DISK.MCR, this particular configuration, this means calling BIOS INT 13h

```
C:\PDOCZ\CHAP4>cmdcopy compile disk c: fat c: 7*70
```

```

25 Abs Disk Read dev 02, at sector 0000001E, 0001 sectors
04 - Input
1302 Read 19 sectors dev 80, head 00, sect 01, trk 06 to 1685:0000
25 Abs Disk Read dev 02, at sector 0000001F, 0001 sectors
04 - Input
1302 Read 19 sectors dev 80, head 01, sect 07, trk 06 to 1705:0000

```

The IN-GENY's plot shows that INT 25h calls device driver function 4 Read, which in turn calls INT 13h AH=2.

If one is to use a RAM disk, for example, the device driver would certainly not call INT 13h. Instead, it would read and write sectors by accessing memory. The construction of RAM disks is an alternative topic that we will take up further below. For an introduction, see the chapter on RAM disk software. For *Microsoft DOS Device Drivers*, and supplement this with an exploration of the new RAM disk model and related issues, such as the issue of extended memory. For example, on 80286 systems, Microsoft's RAMDRIVE.SYS uses the undocumented LOADALL instruction IBM designed for test purposes to allow programs using STDISK.LDR to use memory ram disks in ESDI mode. It is useful to be able to detect whether a given drive is a RAM disk.

For each cluster, function `get_fat_entry()` uses cluster number and DPR information to figure out which FAT sector is needed. Each sector differs from the one `get_fat_entry()` just used; the function calls `read_sector()` to read the FAT's index into the FAT sector; this is trivial for 16-bit FATs and somewhat for 12-bit FATs. For 12-bit FATs, the code also has to worry about a FAT entry spanning two sectors, which allocates space for two sectors and for 12-bit FATs reads in two sectors.

Obviously, the FAT is a linked list of clusters that threads the pieces of each file together and indicates where space is available. Determining the very first cluster for any specified file is tedious, and one can access everything as if. How then is the starting cluster found? That's done by the directory structure, to which we now turn.



## DOS Directory Structure

Every disk volume (that is, each diskette or drives with removable media) or each partition on that removable media has a root directory, which is the starting point for translating human-oriented file names into system-oriented cluster numbers.

The root directory of a medium follows the FAT and precedes the data storage area. Its size is established when the disk is formatted. Unlike most root directories which are just files, the root directory can never change. A typical root directory size for a 500K floppy's FAT carries for a hard disk is usually larger (typically 512 entries). For a floppy disk, DOS won't accept a BPB calling for more than 240 root directory entries. Neither will it accept values other than 2 for the number of FATs, 1 for the number of reserved sectors, or a sector size other than 512 bytes.

Besides having a few files, the root directory has a more important difference from other root directories: The root directory is not accessible from the FAT. To get from one sector to the root directory (the way you cannot work the FAT) is with root directory entries. Instead, all root directory sectors are contiguous. You can use values from the BPB to compute the number of root directory sectors, as shown in the sample program NAMECLUST (Listing 8-5). See the code in `do_cluster()` to handle the case of `cluster == 0`.

As a result of the D, R, E, N, O, Y structure of DOSNTU2.H (Listing 8-3), each directory's subdirectory, which is in the root or a subdirectory, consists of a 32-byte structure. This structure is plainly visible in a hexadecimal dump of a directory structure. In Figure 8-4, the directory containing only sector 5B4Dh from drive C (this sector is the dumped with the `Dos` command) is located in sector 5B4Dh contained a directory for `C:\WINDOWS\24.HAPP` was seen, and using the `NAMECLUST` program, the source code we present as a few moments. In Figure 8-4, the elements `FILESC` and `FILESEX` are visible within this directory, as are the standard `..` and `..` files for the current and parent subdirectories.

Figure 8-4 was generated on a system running StackIt. The directory structure shown is just data inside a `NTMVCHE000` file on drive D, and is not directly on disk. The root entry appears to be on drive C, to make the point a bit more. The system did not need a copy of disk any hard way, because it block device driver can write up the necessary DOS file system data sector on demand.

**Figure 8-4: A DOS Directory Structure**

```
C:\WINDOWS\CHAPS>nameclust .
C:\WINDOWS\CHAPS >> 1448 (sect 25573, D=00005b4d)

C:\WINDOWS\CHAPS>debug
-L 1000 Z 5b4d 1
-d 1000
7c0d 1000 2e 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 10 00 00 00 00
7c0d 1010 00 00 00 00 00 00 20 85-73 1A A8 05 00 00 00 00 00 00 00 00 00 00 00 00 00
7c0d 1020 2e 2e 20 20 20 00 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 10 00 00 00 00
7c0d 1030 00 00 00 00 00 00 20 85 73 1A 12 00 00 00 00 00 00 00 00 00 00 00 00 00 00
7c0d 1040 46 49 4c 45 53 20 20 20 20 43 20 20 20 20 00 00 00 00 00 00 00 00 00 00 00
7c0d 1050 00 00 00 00 00 00 00 50 70 e7 1A AA 22 3c 19 00 00 00 00 00 00 00 00 00 00
7c0d 1060 46 49 4c 45 53 20 20 20 20 45 58 45 20 00 00 00 00 00 00 00 00 00 00 00 00
7c0d 1070 00 00 00 00 00 00 17 89-72 1A 53 05 52 10 00 00
;

C:\WINDOWS\CHAPS>dir
; ***
*          <DIR>      03-19-93  10:41p
*          <DIR>      03-19-93  10:41p
FILES     C           6460 07-07-93   2.02p
FILES     EXE        7506 03-18-93  11:08p
; ***
```





The program gets the short form, of course. Otherwise Chicago would break every DOS program that had obeyed the rules and allocated as little as 64 bytes in which to receive the current directory string. So developers don't have to worry about Microsoft breaking old programs with long filenames. Furthermore, Microsoft is taking care so that, as much as possible, an end-user sees long filenames even when a program sees the short ones. For example, it intends for the Windows common file dialogs (COMMDFL.DLL) to work this way.

However, programs will eventually want to know full filenames. Getting the long filename instead of the short alias requires a new set of functions. Windows programs will use Win32 APIs, and those such as `GetCurrentDirectory` and `CreateFile`. A DOS program will be able to use a new set of INT 21h AX=71XXh functions, where the subfunction in AL is the same as the old DOS API function number. For example, because the old DOS `GetCurrentDirectory` function is INT 21h AH=47h, the new one that knows about long pathnames is INT 21h AX=7147h, while the other registers are identical to the old call; the buffers provided to it by DOS must be large enough to receive the maximum allowed path. Programs can use the new INT 21h AX=4302h `Get Volume Information` function (see below) to get the complete filename allowed path.

Of course, there is a new 21:7160 equivalent to the old 21:60 (undocumented) `FindNext` function. Also, it isn't clear if Chicago will provide functions to convert between long and short filenames.)

In addition to these new INT 21h AH=71h calls, Chicago also requires some new actions for INT 21h AH=57h (4: get last access date, 5: set last access date). Incidentally, Chicago silently skips over the two OS 2 EA actions for INT 21h AH=57h (see Chapter 4) which the previous version had put in comment marks as "reserved". This is important, because some external reports on Chicago (PC Week, August 23, 1993) claimed that Chicago long filenames might cause some deliberate incompatibility with OS 2. The evidence here points in the exact opposite direction. It appears to be deliberately incompatible! A welcome change from the usual incompatibility, seemingly deliberately, discussed in Chapter 1.

Another new DOS call is INT 21h AH=72h (`FindClose`). Whereas previous versions of DOS supported only `FindFirst` and `FindNext`, Chicago requires `FindClose` because the Win32 programming interface (API) supports multiple simultaneous file finds. `FindFirst` (INT 21h AX=7147h) returns a search handle, which you must pass to `FindNext` (INT 21h AX=7148h) and which you must close with `FindClose`.

Long filenames are not supported in the real-mode DOS boot code portion of Chicago, which runs before DOS486.EXE. This means that real-mode programs such as TSRs and device drivers can't call the new long filename APIs if they run at system startup. Similarly,

if few programs run in so-called "Single Application Mode" also won't be able to call the long filename APIs. To determine if the new APIs are available, a program can check the return values from the new 21:71 functions, or it can call the INT 21h AX=4302h `Get Volume Information` function. This new function returns information on the maximum pathnames, component length, whether the file system is case preserving, and whether the long filename APIs are available. The function also returns the file-system name, which in the preliminary release was "LFAT".

Each disk allocated partition (DOS 5.0 called the directory entry, and the one we're most concerned with at the moment) is the cluster word at offset 1Ah. This is the number for the file's first cluster. Recall from FVLC that once you have the first cluster, the FAT chain gives you access to the

rest of the file. The cluster word in the directory entry is what DOS uses to translate a file name to a location on disk.

NAMCLUST.C uses %N% shows you to use the cluster number in directory entries to translate names to locations. NAMCLUST takes a path name, file or directory entry, in a command line, and displays the file or directory's starting cluster number. For convenience when using DIRBC to read sectors from disk, NAMCLUST also translates the cluster number to a hex sector number. For example:

```
C:\WINDOWS\CHAP8>name ust
C:\WINDOWS ==> 18 (sect 493, 0x00001ed)
C:\WINDOWS\CHAP8>debug
-L 1000 2 01ED 1
```

NAMCLUST.C doesn't illustrate the easiest way to turn a path name into a cluster number. You'll see later on that the SF1 entry for an open file contains the file's starting cluster number. Thus opening a file and then looking for its cluster number in the open file's corresponding SF1 is simpler. You could do the same thing with FCBs. Robert Hornik's *Windows Programming for Novice Chapter 3* shows how to use FCBs to find starting clusters again without walking the FAT. However, the code in NAMCLUST.C uses the actual directory structure on disk, and so better illustrates the workings of the DOS file system.

If it's worth just adding to its command line—for example, NAMCLUST.C—the program walks the entire disk's structure, printing out a hierarchical list of every file and its starting cluster. For example, the following shows C:\BORLAND\INCLUDE\SHELL\OC\KIN\ (and other) files:

```
C:\WINDOWS\CHAP8>namclust c
.
.
BORLANDC ==> 172 (sect 2957, 0x00000b8d)
INCLUDE ==> 175 (sect 3005, 0x00000bbd)
SYS ==> 176 (sect 3021, 0x00000bcd)
LOCKING.H ==> 10797 (sect 172957, 0x0002a59d)
STAT.H ==> 10798 (sect 172973, 0x0002a5ad)
TIMED.H ==> 10799 (sect 172989, 0x0002a3bd)
TYPES.H ==> 10800 (sect 173005, 0x0002a3cd)
ALLOC.H ==> 10711 (sect 171581, 0x00029e3d)
ASSERT.H ==> 10712 (sect 171597, 0x00029e4d)
BCD.H ==> 10713 (sect 171613, 0x00029e5d)
BIOS.H ==> 10715 (sect 171645, 0x00029e7d)
.
etc
```

In NAMCLUST.C, the variable `local` controls this listing of the entire disk.

## Listing 8.5. NAMCLUST.C

```
/*
NAMCLUST.C -- Convert file name to starting cluster
Andrew Schulman, July 1993
Overall structure
main
    truename -- to get canonical filename
    get_dpb
    do_cluster -- process a directory cluster, starts with root dir <+
        read_sectors (if it's a root directory)
        read_cluster
        read_sectors (INT 25h)
    do_dir
        printf -- success! print results
        do_cluster -- partial match, recurse to next level
        get_fat_entry -- get next cluster of directory
    fail -- no such file!
*/
#include <stdlib.h>
```

```

#include <stdio.h>
#include <string.h>
#include <ctype.h>
#include <dos.h>
#include diskstuf.h"

// see Nameless, Assembly Language Lab Notes, p. 88
#define CLUSTER_TO_LSNH(cluster) \
    (((DWORD) (cluster) - 2) * sect_per_clust) + first_sect)

void do_cluster (WORD cluster),
void do_dir (WORD cluster, DIR_ENTRY *buffer),
void read_cluster (int drive, BYTE far *buffer, WORD cluster),

static STRING glob_name[64] = {0},
static DPB far *glob_dpb,
static WORD want_cluster = 0xffff,
static WORD bytes_per_clust = 0, sect_per_clust = 0, first_sect = 0,
static int leve = 0, glob_num_dir = 0, glob_drive,
static int num_levels = 0;
static char canon[128];
static int do_all = 0;

void fail (const char *s) { puts(s); exit(1); }

main (int argc, char *argv[])
{
    char partname[64],
    char *s, *s2,
    int i, part;

    if (argc < 2)
        fail ("usage: nameclust [pathname or drive:~];");

    for (i=0; i<64; i++)
        glob_name[i] = (STRING) malloc(16);

    // program uses DOS 3.31 style of INT 25h
    if (! _osmajor < 3) || !_osminor || !_osminor < 31)
        fail ("this program requires DOS 3.31 or higher");

    if (! strlen(argv[1]) < 2 && argv[1][1] == ':')
    {
        do_all++;
        num_levels = 0xffff;
        strcpy (canon, strdup(argv[1]));
    }
    else
    {
        if (! truenam(argv[1], canon))
            fail ("can't get truenam"); // might be JOINed drive
        if (canon[1] != ':')
            fail ("Sorry, this doesn't work on network-redirector drives");

        for (part=0, s=canon, s2=partname; s++)
            if ((*s == '\\0') || (*s == '\\'))
            {
                *s2 = '\\0';
                strcpy(glob_name[part], partname);
                part++;
                if (*s == '\\0')
                    break;
                s2 = partname,
            }
        else
            *s2++ = *s;

        num_levels = part - 1,
        level = 1;
    }
}

```

```

glob_drive = canon[0] - 'A';
if (! (glob_dpb = get_dpb(glob_drive+1)))
    fail("can't get DPB");

sect_per_clust = glob_dpb->sectors_per_cluster + 1,
bytes_per_clust = glob_dpb->bytes_per_sect * sect_per_clust;
first_sect = glob_dpb->first_data_sector;
glob_num_dir = bytes_per_clust / sizeof(DIR_ENTRY);
do_cluster(0); // kick-start search
}

void do_cluster(WORD cluster)
{
    BYTE *buffer,
    if (cluster == 0) // phony cluster # for root directory
    {
        int num_sect,
        buffer = (BYTE *) calloc(glob_dpb->max_root_dir, sizeof(DIR_ENTRY));
        if (! buffer)
            fail("insufficient memory");
        num_sect = (glob_dpb->max_root_dir * sizeof(DIR_ENTRY)) /
            glob_dpb->bytes_per_sect;
        read_sectors(glob_drive, buffer, num_sect,
            (_osmajor >= 4) * glob_dpb->vers dos4 first_dir_sector :
            glob_dpb->vers dos3.first_dir_sector);
    }
    else
    {
        if (! (buffer = (BYTE *) calloc(bytes_per_clust, 1)))
            fail("insufficient memory");
        read_cluster(glob_drive, buffer, cluster);
    }
    do_dir(cluster, (DIR_ENTRY *) buffer);
    // shouldn't get here: can't free!
    if (do_all)
        free(buffer);
    else
        fail("can't find file"); // error if get here
}

void do_dir(WORD cluster, DIR_ENTRY *buffer)
{
    DIR_ENTRY *dir, *p,
    char name[16], *s;
    WORD next, num_dir,
    int entry, i;

    dir = (DIR_ENTRY *) buffer,
    num_dir = (cluster == 0) * glob_dpb->max_root_dir : glob_num_dir,
    for (entry=0, p=dir, entry<num_dir, p++, entry++)
    {
        if (p->name[0] == '\0') return; // end of dir
        if (p->name[0] == 0xE5) continue; // deleted entry
        if (p->name[0] == ' ') continue; // don't bother with . and .
        // looks like Chicago uses these for long-filename entries
        if (p->attrib == 0x0F) continue;
        memcpy(name, p->name, 8); name[8] = 0;
        s = name, while (*s && *s != ' ') s++; *s = '\0';
        if (name[0] == 0x05) name[0] = 0xE5, // 0xE5 is valid first char
            if (p->ext[0] != ' ')
            {
                strcat(name, ".");
                strcat(name, p->ext, 3);
                s = name; while (*s && *s != ' ') s++, *s = '\0';
            }
    }
}

```

```

if (do_all || strcmp(name, glob_name[level]) == 0) // got match'
{
    if (do_all || (level == num_levels) // done
    {
        DWORD sector = CLUSTER_TO_LSN(p->cluster);
        char *s;
        if (do_all)
        {
            int i;
            for (i=0; i<level; i++)
                printf(" ");
            s = name;
        }
        else
            s = canon;
        printf("%s ==> %u (sect %lu, D%LOB(s))\n",
            s, p->cluster, sector, sector);
        if (!do_all)
            exit(0);
    }
    if (p->attrib & DIRECTORY_ATTR) // already know not or ..
    {
        level++;
        do_cluster(p->cluster); // if directory, recurse
        level--; // for do_all
    }
}

// create next directory cluster
if (linux_get_fat_entry(glob_drive, glob_dpb, cluster) < DxFFFD)
    do_cluster(next);
}

void read_cluster(int drive, BYTE far *buffer, WORD cluster)
{
    DWORD sector;
    if (cluster == 0) sector = CLUSTER_TO_LSN(cluster);
    else if (strcmp("A:", glob_drive) == 0) sector = glob_dpb->vers dos3 first_dir_sector;
    else sector = glob_dpb->vers dos3 first_dir_sector;
    if (!read_sector(drive, buffer, sect_per_clust, sector))
        fail("can't read cluster");
}

```

**NAME** USES `get_fat_entry` to get disk's directory hierarchy, trying to match the name of disk specified file or directory.

The program calls the `truename()` function (INT 21h AH=60h in DISKSUBC (Listing B-4)) to canonicalize the name specified on its command line (see "Finding a True Name").

## Finding a True Name

In the MS-DOS file system, things may not be what they seem. A file called `C:\FLOPPY\FOO\BAR` may actually be located on a joined floppy disk in drive A:, and a subdirectory called `F:\SOURCEFS` may actually be located on a network file server (probably not even running MS-DOS) in a directory called `BIN\EXPORT\DOS`. A canonical (true) path string resolves all these log cat (that is, non-physical) drive and path references to an absolute pathname, taking account of any renaming due to JOIN, SUBST, or network redirections.

Fortunately, there is a DOS function that provides a true canonical pathname: undocumented INT 21h function 60h (Resolve Path String to Canonical Path String). This corre-



sponds to the undocumented TRUENAME command in COMMAND.COM in DOS 4.01 and higher (see Chapter 10). For example:

```
C:\UNDOC2>truename foo.bar
C:\UNDOC2\FOO.BAR
C:\UNDOC2>subst g: c:\undoc2
C:\UNDOC2>truename g:\foo.bar
C:\UNDOC2\FOO.BAR
C:\UNDOC2>c:\dos\join a: c:\floppy
C:\UNDOC2>truename c:\floppy\foo.bar
A:\FOO.BAR
C:\UNDOC2>rem E: and F: are network drives
C:\UNDOC2>truename f:\foo.bar
\\HOME\LAN\ANDREW\FOO.BAR
C:\UNDOC2>truename g:*.* exe
\\BIN\EXPORT.DOS\?????.???.EXE
```

Because G: is a SUBSTed alias for C:\UNDOC2, the true name for C:\FOO.BAR is C:\UNDOC2\FOO.BAR. Likewise, because C:\FLOPPY is a JOINed alias for the A: drive, A:\FOO.BAR is the true name for C:\FLOPPY\FOO.BAR. F: is located on a Sun SPARCstation (made available to DOS using PC TCP from FTP Software), and the \ that starts off the truename of F:\FOO.BAR is the universal naming convention (UNC) indicating that this is a network drive.

It is important to realize, however, that the file FOO.BAR need not necessarily exist. This is a source of tremendous confusion about what function 60h and the TRUENAME command do. The truename function and TRUENAME command deal with path names, not with actual files. This is important for disk utilities such as NAMCUST because it provides an easy way of determining that the user has specified a device with which the program can work. Simply examining the specified drive letter such as C: or A: tells you nothing about the actual drive involved. It is also useful sometimes to have some assistance from the operating system in interpreting complex pathnames with many \ and subdirectories.

```
C:\UNDOC2>truename foo\bar\...\ \ \ \
C:\UNDOC2
```

The undocumented TRUENAME command in COMMAND.COM in DOS 4.01 and higher is simply a wrapper around undocumented function 60h, which became available in DOS 3.0. You can call function 60h using the truename() function from DISKSTUFF.C (Listing 8-4). Like most of the code in this chapter, the truename() function in DISKSTUFF.C is only callable from a real mode DOS program. To call INT 21h AH: 60h from a protected mode Windows program, you need to use the techniques from Chapter 3 of this book; in particular, Listing 3-20 shows a protected mode version of truename() that uses DPMI. This is important because the normally well-informed *Microsoft Systems Journal* claimed (July-August 1992) that, unfortunately, Truename cannot be called from Windows "Can too!"

One problem with function 60h is its slightly odd interaction with JOIN. If you JOIN A:\C:\FLOPPY, then while TRUENAME C:\FLOPPY\FOO.BAR works, TRUENAME A:\FOO.BAR gets an "invalid drive specification" message! Meanwhile, function 60h does seem to work

properly with the odd ASSIGN command. For example, if you ASSIGN a=c, then TRUENAME A FOO BAR properly returns C:\FOO BAR.

Another potential problem is that function 60h hits the disk. In some situations, accessing the Current Directory Structure might be preferable to function 60h. With the exception of its coexistence with ASSIGN, function 60h relates fairly directly to the CDS. For any drive n, the output of function 60h for the string "n" (that is, for the subdirectory) matches cds[drive n] - current\_path. This point will make more sense when we discuss the CDS in detail, later in this chapter.

TRUENAME is a small sample program on the accompanying disk that runs truename() in a loop over the current directory for each drive on your system. The result is similar to that of the ENUMDRV program, shown later in Listing 8.11. Apparently there may be problems calling INT 21h AH 60h repeatedly under NetWare.

**After calling truename:** NAMCLUSI then calls get\_dpb() to retrieve the DPB for the specified disk. NAMCLUSI uses the DPB to determine values such as the number of directory entries in a directory cluster and the sector of the sector for the root directory. NAMCLUSI then calls do\_cluster() to kick-start the name search in the root directory. For each directory do\_cluster() reads in the entire root directory. Root directories have a fixed number of entries which is given in the DPB; you don't need to worry about the FAT. For non-root directories do\_cluster() reads a single directory cluster into memory and then calls do\_dir() to process the directory cluster.

Finally, do\_dir() examines every directory entry (since it the name matches the current portion of the path name) to see if the specified path name is found at the end of the specified path name. If not, it calls do\_dir() again. If the path name remains to be matched, do\_dir() moves into subdirectories by calling do\_cluster(). Finally, do\_dir() hits the end of a directory and either matches it successfully or fails. To determine if there are additional clusters for the directory, do\_dir() reads do\_dir() to read them in. If NAMCLUSI arrives at the bottom of the case, without coming to the program has failed to match the specified path name. Therefore, it matches path name exists.

Thus, the task of NAMCLUSI shows how to turn pathnames into starting cluster numbers. For example, we could boot this code onto the front of FAT12 or FAT16 to produce a more useful program that takes a path name, rather than a cluster number, as input and lists the file or directory's possible fragmented FAT chain.

It is also possible to work in the opposite direction. Given a starting cluster number, this program could also convert a corresponding file path name. This will be useful later when this chapter examines the DOS System File Table (SFT). The SFT contains only the name and extension of a file, without any path information. However, the SFT does contain the starting cluster number for any open file, and some programs do figure out the full path name. Unfortunately, this CHUSINAM operation can take a long time. This script by Norton Utilities SA program has an "Information on item" option, choosing first to access cluster number causes SFT to put up "Working one moment please..." notice and grind away for a while on the disk. Potentially every subdirectory on the disk must be examined to match a cluster number.

It is also possible to work in the opposite of NAMCLUSI. In Listing 8.5, takes a drive and cluster number on its command line, and walks the directory structure, trying to produce the corresponding path name.

```
C:\UNDOC2\CHAP8>clustnam d: 28005
C:\UNDOC2\CHAP8\CHAP8.TXT
```

Running NAMCLUSI confirms this

```
C:\WINDOWS\CHAP8>namclust chap8.txt
C:\WINDOWS\CHAP8 TXT ==> 28005 (sect 112289, 0x0001b6e1)
```

The `do_cluster` and `do_dir` functions in `CLUSTNAM.C` differ those from in `NAMCLUST.C`. Rather than creating two slightly different versions of the same basic code, we should have written a more general, sophisticated library for manipulating DOS directory entries.

### Listing 8-6: CLUSTNAM.C

```
/*
CLUSTNAM.C -- Convert cluster # to full pathname
See also NAMCLUST.C (Convert pathname to cluster #)
Andrew Schulman, July 1993

Program needs large model: bcc -ml clustnam.c diskstuf.c

Overall structure
main
  get_dpb
  search_for_cluster
    do_cluster <-----*
      read_sectors |
      read_cluster |
      read_sectors |
    do_dir |
      do_cluster -- *
      get_fat_entry
*/

#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <ctype.h>
#include <dos.h>
#include "diskstuf.h"

void search_for_cluster();
void do_cluster(WORD cluster);
void do_dir(WORD cluster, DIR_ENTRY *buffer);
void read_cluster(int drive, BYTE far *buffer, WORD cluster);
void fail(const char *s) { puts(s); exit(1); }

static STRING glob_name[64] = {0};
static DPB far *glob_dpb;
static WORD want_cluster = 0xffff;
static WORD bytes_per_cluster = 0, sect_per_cluster = 0, first_sect = 0,
static int leve = 0, glob_num_dir = 0, glob_drive;

main(int argc, char *argv[])
{
    int i;
    if (argc < 3)
        fail("usage: clusters [drive] [cluster]");
    // program uses DOS 3.31 style of INT 25h
    if (!osmajor < 3) || !osmajor == 3 && osminor < 31)
        fail("this program requires DOS 3.31 or higher");
    for (i=0; i<64; i++)
        glob_name[i] = (STRING) malloc(16);
    glob_drive = toupper(argv[1][0]) - 'A';
    want_cluster = atoi(argv[2]);
    if (glob_dpb = get_dpb(glob_drive+1))
    {
        if (want_cluster <= glob_dpb->highest_cluster)
        {
            search_for_cluster();
        }
    }
}
```

```

        fail("couldn't find cluster"), // error if get here
    }
    else
        fail("cluster number too high");
}
else
    fail("can't get DPB"),
    return 0;
}

void search_for_cluster(void)
{
    glob_drive = glob_dpb->drive;
    sect_per_clust = glob_dpb->sectors_per_cluster + 1;
    bytes_per_clust = glob_dpb->bytes_per_sect * sect_per_clust;
    first_sect = glob_dpb->first_data_sector;
    glob_num_dir = bytes_per_clust / sizeof(DIR_ENTRY);
    do_cluster(0), // phony cluster 0 to kick start
}

void do_cluster(WORD cluster)
{
    BYTE *buffer;
    if (cluster == 0) // phony cluster # for root directory
    {
        int num_sect;
        buffer = (BYTE *) calloc(glob_dpb->max_root_dir, sizeof(DIR_ENTRY));
        if (! buffer)
            fail("insufficient memory");
        num_sect = (glob_dpb->max_root_dir * sizeof(DIR_ENTRY)) /
            glob_dpb->bytes_per_sect;
        read_sectors(glob_drive, buffer, num_sect,
            [ osmajor >= 4 ] ? glob_dpb->vers.dos4.first_dir_sector :
                glob_dpb->vers.dos3.first_dir_sector);
    }
    else
    {
        if (! (buffer = (BYTE *) calloc(bytes_per_clust, 1)))
        {
            if (sizeof(void *) != 4)
                puts("Hey, you shru'd have compiled with large model!");
            fail("insufficient memory");
        }
        read_cluster(glob_drive, buffer, cluster);
        do_dir(cluster, (DIR_ENTRY *) buffer);
        free(buffer);
    }
}

void do_dir(WORD cluster, DIR_ENTRY *buffer)
{
    DIR_ENTRY *dir, *p;
    char name[9], ext[4];
    WORD next, num_dir;
    int entry, i;

    dir = (DIR_ENTRY *) buffer;
    num_dir = (cluster == 0) ? glob_dpb->max_root_dir : glob_num_dir;
    for (entry=0, p=dir, entry<num_dir, p++, entry++)
    {
        if (p->name[0] == '\0') return, // end of dir
        if (p->name[0] == DxE5) continue, // deleted entry
        if (p->name[0] == '.') continue, // don't bother with . and ..

        // looks like Chicago uses these for long-filename entries
        if (p->attrib == 0x0F) continue;
    }
}

```

```

memcpy(name, p->name, 8); name[8] = 0;
if (name[0] == 0x05) name[0] = 0xE5; // 0xE5 is valid first char
memcpy(ext, p->ext, 3); ext[3] = 0;
strcpy(glob_name[level], name);
if (ext[0] != ' ' || 0)
{
    strcat(glob_name[level], ".");
    strcat(glob_name[level], ext);
}

if (p->cluster == want_cluster) // found it!
{
    char *s;
    int i;
    putchar(glob_drive + 'A'); putchar(':'); putchar('\');
    for (i=0; i<level; i++)
    {
        s = glob_name[i]; // trim spaces out of name
        while (*s) { if (*s != ' ') putchar(*s); s++; }
        putchar('\');
    }
    s = glob_name[level]; // trim spaces out of name
    while (*s) { if (*s != ' ') putchar(*s); s++; }
    putchar('\n');
    exit(0); // done: success!
}

if (p->attrib & DIRECTORY_ATTR) // already know not dir
{
    level++;
    do_cluster(p->cluster); // if directory, recurse
    level--;
}

// locate next directory cluster
next = get_fat_entry(glob_drive, glob_dpb, cluster);
if (next < 0xFFFD) // end of dir
    do_cluster(next);
}

// see Nummel, Assembly Language Lab Notes, p 88
#define CLUSTER_TO_LSN(cluster) \
(((DWORD) (cluster) * 2) * sect_per_cluster) + first_sect

void read_cluster(int drive, BYTE far *buffer, WORD cluster)
{
    DWORD sector;
    if (cluster != 0) sector = CLUSTER_TO_LSN(cluster);
    else if (_osmajor >= 4) sector = glob_dpb->vers dos4 first_dir_sector;
    else sector = glob_dpb->vers dos3 first_dir_sector;
    if (! read_sectors(drive, buffer, sect_per_cluster, sector))
        fail("can't read cluster");
}

```

The `NAMCLUSTER_CLUSTERNAME` uses the `DISKUFF` module from Listings 8-2 and 8-4. The program's overall structure is very similar to that of `NAMRFILE`, except that where `NAMRFILE` starts out with the directory search as soon as it finds a mismatch, `CLUSTERNAME` keeps walking the entire disk directory hierarchy until it finds a match.

### The Drive Parameter Block (DPB)

The programs presented so far in this chapter all relied heavily on the `DPB` structure and `get_dpb` function in `DISKUFF` ( Listings 8-3 and 8-4) to determine characteristics of a drive, such as its

bytes per sector, sectors per cluster, sectors per FAT, location of the first FAT, directory, data sectors, and so on. But we haven't said much yet about the DPB.

For each disk device (if assigned), there is a Drive Parameter Block. These 32-byte blocks contain information that DOS uses to convert cluster numbers into logical sector numbers for passing to the disk device driver. Read and Write functions and to associate the device driver for that device with its assigned drive letter.

DOS uses the DPB for each drive immediately after calling the driver's Initialize routine during the boot process. Back in Chapter 7, you saw the `INITDISK` program, which mimics DOS operation (and block device drivers). DPB for the drivers is `INITDISK` or `INITBIOS.COM` (normally floppy `A:` and `B:` together with an `hdisk.C` associated with `IO.SYS`) initializes itself before processing `CONTCFG.SYS`. For other block devices, they are created as one of the final steps of installing the device driver, while processing `CONTCFG.SYS`.

As you saw in Chapter 7, `INITDISK` uses the still undocumented `INT 21h` function 33h (passing the `DP` pointer to the driver's DPB; see the `PrintDiskDev` function in `DEVINFO.C` in Chapter 7, Listing 7-78). You should recognize many of the fields in the DPB from our earlier discussions of the DOS files; also the BIOS description of the value. A copy of the DPB would be useful to the device developer, and the pointer is part of the information returned by the driver's Initialize routine, device driver function 0. As you saw earlier, for the built-in disk drives, the DPB is maintained in the boot sector of each volume, and each time a volume is changed, DOS uses its DPB to rebuild the DPB using the new volume's characteristics (differ from the original values).

Not all DPBs are for disk drives (which have an `hdisk.C` associated with them); the exception to this is `drive E:` in a single floppy's case; it's always assumed to exist, even when it doesn't. On the other side, most non-disk devices (but non-serial and drives, such as RAM disks, generally do have DPBs).

Each DPB is added to the rest of the FAT pointer to the DPB structure. `FFFFh` in the pointer's offset position indicates the end of this linked list. But which DPBs are chained together in a linked list? A user can't say that; since the `IO.SYS` or `BIOS` is a static save copy, the DPB for a given drive is given by DOS via the DPB pointer `INT 21h` function 32b (which we saw in Chapter 7 as part of `DEVINFO.C`), or for a single block device, and which is often used at disk programs such as Norton Utilities or `FDISK`. Microsoft has barely documented this function in its *MS DOS Programmer's Reference*. Even so, according to Microsoft, this function is only for MS-DOS 5.0 and higher. (See `fdisk.com`.) In Microsoft documentation, you can't write disk utilities that work under any previous version of DOS.

The `fdisk.com` program in Listing 5-7, `DPBINFO.C`, uses `INT 21h` function 32b to display capacity information for each drive in the system with a DPB. For example, on a system with a 1.2 megabyte floppy drive, a 400K floppy, and the tank, a 0.8 megabyte hard disk, and a 64K RAM drive installed with `INITDISK`, it gives the following as the program's output:

```
Drive A: 512 bytes/sector * 2 sectors/cluster =
          1024 bytes/cluster * 354 clusters = 362496 bytes
Drive C: 512 bytes/sector * 8 sectors/cluster =
          4096 bytes/cluster * 17648 clusters = 72786208 bytes
Drive E: 512 bytes/sector * 1 sectors/cluster =
          512 bytes/cluster * 122 clusters = 62464 bytes
```

The program displays this information twice: once by walking the DPB linked list (whose head is at `fsd[0]` in `IO.SYS`), and once by calling `INT 21h` function 32b (to get `dpb` from `DISKINFO` in Listing 8-4) for each drive's last drive.

A lot of interest (which you'll work later) that walking the DPB chain for `MS-DOS` drives, for example, is to associate the DPB with the drive's Current Directory Structure. For an illustration, see the function `mapdrive` in `hdisk.C` in `INITDISK.C` (Listing 8-11) which can retrieve a DPB with the expression `currentdrive > dpb`.

## Listing 8-7: DPBTEST.C

```

/*
DPBTEST.C -- uses undocumented INT 21h function 32h (Get DPB)
to display bytes per drive, but first walks the DPB chain,
showing the difference between the two access methods.
A third access method gets the DPB from the CDS (see
maybe_ram_disk() in ENUMDRV.C).
Andrew Schulman, updated July 1993
*/

#include <stdlib.h>
#include <stdio.h>
#include <dos.h>
#ifdef __TURBOC__
#include <dir.h>
#endif
#include "diskstuf.h" // DPB structure, get_dpb()

void fail(const char *s) { puts(s), exit(1); }
void display(DPB far *dpb)
{
    DWORD bytes_per_clust =
        dpb->bytes_per_sect * (dpb->sectors_per_cluster + 1);
    printf("Drive %c: %A * %d = %d\n",
        dpb->drive, dpb->bytes_per_sect,
        dpb->sectors_per_cluster + 1,
        bytes_per_clust);
    printf("clusters = %d, highest cluster = %d\n",
        dpb->highest_cluster - 1,
        dpb->highest_cluster - 1);
}

main()
{
    BYTE far *sysvars;
    DPB far *dpb;
    unsigned lastdrive;
    int i;

    if ((_osmajor < 3) || (_osmajor == 3 && _osminor < 2))
        fail("This program requires DOS 3.2 or higher");

    puts("Using DPB Linked List");
    _asm mov ah, 52h // get SysVars *
    _asm int 21h
    _asm mov word ptr sysvars+2, es
    _asm mov word ptr sysvars, bx
    /* pointer to first DPB at offset 0h in SysVars */
    if (! (dpb = *((DPB far * far *) sysvars)))
        return 1;

    do {
        // don't show both A: and B: etc.
        if ((! _dos_driverremoveable(dpb->drive + 1)) ||
            (_dos_getdriveemp(dpb->drive + 1) == dpb->drive + 1))
            display dpb;
        dpb = (_osmajor < 4) ? dpb->vers dos3.next : dpb = dpb->vers.dos4.next;
    } while ((FP_OFF(dpb) != 0xFFFF);

    // Another (better) method gets DPB pointer out of drive's CDS;
    // #define GET_DPB(drv) (currdir(drv)->dpb)
    // see maybe_ram_disk() #ifdef USE_CDS_DPB in Listing 8-11

    puts("Using INT 21h function 32h");
    _dos_setdrive(0xFF, &lastdrive), // get lastdrive
    for (i=1; i<lastdrive; i++)

```

```
if (dph = get_dph()) // checks for removable, critical error
    display(dph),
```

```
return 0,
```

This program brings up an important reason to use INT 21h function 32h instead of walking the DPB linked list. For removable media, function 32h goes to the disk and therefore picks up the most current information. Walking the disk list merely gets whatever possible state DPB happens to be in memory. Even access to 360k floppy disk in drive A, remove it, put in a 1.2 megabyte floppy with out accessing it, and then walk the DPB linked list, you get the DPB for the 360k floppy. Function 32h would not make this mistake.

On the other hand, for non-386 bits the disk see the disassembly of functions 32h and 11h in Chapter 6. Also, makes it convenient to get the DPB of drives with removable media. Further, you want to avoid doing two drives A and B in a system where these logical drives are mapped to the same physical floppy drive.

The version of DPB INT 15 listing 8-7 differs substantially from that in the first edition of *Undocumented DOS*, which had tried to deal with the above problem by checking for drives A and B, missing the fact that other variables such as those created with DRIVER.SYS may involve removable media. This error is grave, was further compounded by changing the floppy disk logical drive indicator at address 504h. Totally hopeless!

In real situations, should one generic IO-PL calls to determine whether a device uses removable media, INT 21, AX=4080, and 32h for the logical drive map, INT 21, AX=4401h. These calls, along with an INT 24h critical error handler, make it clear there is no media in the drive at all, are incorporated into the get\_dphs function in the DPBUTILS section of DSKUTIL.C listing 8-4.

One strange aspect of DPB. Many critical DOS disk errors were thrown into temporary confusion in the mid-1980s. DOS 4.0 was set of some bugs, change to the DPB structure, the sectors per FAT field, a sector 0's use, the appendix grew from a *how-to* to a *how-to* so all subsequent fields were supplied, a sector 0's use, a noted at the time, and Mirrored. Functions 52h in DOS 7.0 *Lab Jour* and the year 1989, may be a few modifications probably a major impact in the Norton Utilities and their successors that relied on the undocumented DOS data structure. Rather than becoming increasingly complex, this sort of change is a good excuse to let up customers with an upgrade release.

## Buffers and Disk Caches

Like any proper operating system, MS-DOS has buffers. I/O. Rather than directly reading sectors off the disk, DOS first checks to see whether the sector is already present in an in-memory buffer. DOS uses a byte table, the buffer table, to record data sectors. DOS buffers sectors, not higher-level clusters or other file objects. The PUBLICS statement in CONFIG.SYS controls the number of sector buffers, which may be located in or just recently used. FRL, a regular linked list, SysVars holds a pointer to the head of this list.

Each buffer sector buffer contains a header which identifies the drive, currently using that buffer, the sector or cluster, the cylinder, starting byte, indicating what type of sector (FAT, directory, or data) it contains. It points to the next buffer header in the chain.

The buffer table was not added to DOS 2.0. In DOS 1.x, there was a single sector buffer, Jim Paterson admitted this was "a design mistake, made as their system could not defend." On the other hand, while DOS 1.x kept the FAT in memory, could not do this, marked contrast to CP/M which could require a floppy disk reads just to find the location of a user's data. DOS 2.0 and higher rely entirely on the buffers, or keeping often used FAT sectors in memory. Paterson notes "An Inside Look at MS-DOS," *Byte*, June 1983.



The new MS-DOS does not keep the file allocation tables in memory at all times. Instead, file tables share the use of the sector buffers. This means that at any one time, all part or none of FAT may be in memory. The buffer handling algorithms will presumably keep some used sectors in memory, and this applies to only dual sectors of the FAT as well. This change in the DOS goes completely against my original design principles. Now we're back to doing disk reads just to find out where the data is.

With today's regime, a memory resident FAT could occupy up to 128k/64k clusters \* 16 bits) of memory.

DOS uses the buffers in sequence, changing the linkages as necessary to maintain the most recently used buffers near the front of the chain. Any DOS sector access first walks through the chain of buffers, looking for the expected drive and sector; if found, it can use the buffer contents without having to hit the disk. Moving a new used buffer up to the front of the chain guarantees that a new search reaches the end of the chain without finding its sector; the buffer at the end would be the last recently used and, thus, the proper one, in its scheme. It replaces with the new data read from the disk.

Unfortunately, this simple approach did not take into account the pattern by which DOS performs disk reads. In practice, the buffer chain filled rapidly with FAT and directory data, leaving only a few buffers for all file data. Careless file systems was modified several times under DOS 2.0 and 3.0, but performance problems remained significant. The DOS buffers underwent a major implementation change in DOS 4.0 when IBM introduced a complicated flushing scheme and a mechanism for recycling buffers or expanded memory. This was thrown out in DOS 5.0, which returned to the simpler LRU buffers scheme.

In DOS 3.11 and DOS 5.0 and higher keep the buffers in the HMA. Also for the first time in DOS 5.0, sectors accessed with the INT 25h and INT 26h absolute disk read and write functions first check the sector buffers.

Instead of a pointer to the head of the buffers chain, SysVars in DOS 4.0 and higher also contain the *number* of buffers that is the value from the BUFFERS statement in CONFIG.SYS (or more of them - see SysVars or PC's List of Lists" in this chapter). Determining BUFFERS is probably the only practical way a program could use buffers information. It is difficult when running a program to determine which CONFIG.SYS file was used to boot the system. In fact, prior to the advent of the INT 26h function, even in DOS 4.0, one couldn't even tell what drive the sector was accessed from. So, a shell script and a signature program is all you want to determine the value of BUFFERS and FILES (the next section).

To know the value of BUFFERS in earlier versions of DOS, a program must walk the buffers chain as shown in BUFFERS.C (Listing 8.8). In DOS 4.0 and higher, this program also prints out a description of each buffer's contents. For example:

```
C:\MSDDC7\CHAP8>buffers
FFFFA8C8 -- C: #208 -- DIR
FFFFB96B C #3135 DATA
FFFFA4A0 C #7 FAT
FFFF9A3C C #3226 -- DATA
FFFFB540 D #109 FAT
FFFFB090 D #110 FAT
etc.
FFFFB11B -- A #19 DIR
FFFF9400 -- A #20 -- DIR
BUFFERS=30
```

As you can see, the buffers include floppy diskettes as well as hard disks, however, network redirected drives are not included. It is instructive to run BUFFERS, then perform a disk operation, such as DIR, or run some other program, and then run BUFFERS again. You can see the contents of the buffers change. Of course, running BUFFERS EXE itself changes the contents of the buffers.

When DOS's HIGH-DOS usually allocates the buffers in the high memory area (HMA) just above 1 MB, this area can be accessed only if the processor's A20 address line is enabled. If the buffers are in the HMA but A20 is off, `BUFFERS.C` in Listing 8-8 calls `VMS` to enable A20. This code is very similar to that MS-DOS used back when DOS-HIGH but A20 is off (see Chapter 6). Other programs in this chapter that access structures such as the CDS and SEI don't contain code to enable A20 because DOS's late doesn't allocate these structures in the HMA. Some third-party structures do rely on the CDS and SEI to upper memory (UMBs), but A20 is irrelevant to UMB access.

Low programs don't yet access the DOS buffers, and so moving them to the HMA "breaks" few if any programs. I'm not sure moving the CDS or SEI to the HMA might have disastrous results for programs that didn't know to check A20 first. This might become important because Chicago is required to have CONFIG.SYS settings such as `LASTDRIVE=HIGH` and `FILES=HIGH`, though these probably use UMBs rather than the HMA. On the other hand, any problems will be quite rare, because as we saw in Chapter 6, MS-DOS leaves A20 on, and because any initial DOS call will put A20 back on if some program turns it off.

### Listing 8-8: `BUFFERS.C`

```
/*
BUFFERS.C   Display buffer chain, and count BUFFERS-
See also COUNTF.C to determine value of FILES-
Andrew Schulman, revised July 1993
Added code to check and enable A20 if BUFFERS are in HMA
*/

#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <dos.h>

typedef unsigned char BYTE;
typedef unsigned short WORD;
typedef unsigned long DWORD;

#define MK_FP
#define MK_FP(seg, ofs) \
    ((void Far *) ((DWORD) (seg) << 16) | (ofs))
#endif

#pragma pack(1)
typedef struct dskbuffs {
    struct dskbuffs far *next;
    BYTE drive, flags,
    } DSKBUFFS;

typedef struct {
    WORD next, prev,
    BYTE drive, flags,
    DWORD sector,
    } DSKBUF4;

#pragma pack(1)

void fail(const char *s) { puts(s), exit(1); }
char *buff_status(BYTE flags)
{
    static char buf[128],
    buf[0] = '\0';
    if (flags & 2) strcpy(buf, "FAT ");
    if (flags & 4) strcat(buf, "DIR ");
    if (flags & 8) strcat(buf, "DATA ");
    if (flags & 16) strcat(buf, "REF "); // referenced
    if (flags & 32) strcat(buf, "DIRTY "); // modified, not yet written
    if (flags & 64) strcat(buf, "REMOTE ");
    return buf;
}
}
```

```

/* XMS stuff in case BUFFERS in HMA *****
/* If XMS present, store entry point in xms_entrpoint */
static void (far *xms_entrpoint)(void) = (void (far *)()) 0;
int xms_is_present(void)
{
    int available = 0;
    if (xms_entrpoint) return 1;
    _asm {
        mov     ax, 4300h
        int     2fh
        cmp     al, 80h
        je      present
        jmp     SHORT done
    }
present:
    available = 1,
    _asm {
        mov     ax, 4310h;
        int     2fh;
        push   ds;
        mov     ax, seg xms_entrpoint;
        mov     ds, ax;
        mov     word ptr xms_entrpoint, bx;
        mov     word ptr xms_entrpoint+2, ax;
        pop    ds;
    }
done:
    return available;
}

int xms_local_a20(int func)
{
    if (! xms_entrpoint)
        if (! xms_is_present())
            return 0; // no XMS
    _asm mov ah, byte ptr func
    (*xms_entrpoint)();
    // return value in AX; success=1
}

int xms_local_enable_a20(void) { return xms_local_a20(5); }
int xms_local_disable_a20(void) { return xms_local_a20(6); }
static int testing = 0,
/*****
unsigned buffers(void)
{
    BYTE far *doslist,
    WORD buffers;

    _asm mov ah, 52h
    _asm int 2fh
    _asm mov word ptr doslist+2, ax
    _asm mov word ptr doslist, bx
    /* pointer to first disk buffer in list of Lists */
    if ( _osmajor < 3)
        /*! This program requires DOS 3.0 or higher!;
    else if ( _osmajor == 3)
    {
        DSKBUF3 far *dskbuf3,
        if ( _osminor == 0)
            dskbuf3 = *((DSKBUF3 far * far *) (doslist + 0x13)),
        else

```

```

    dskbuf3 = *((DSKBUF3 far * far *) (doslist + 0x12));
    for (buffers=1, ; buffers++)
    {
        printf("buffer @ %p\n", dskbuf3);
        if ((dskbuf3 = dskbuf3->next) == (DSKBUF3 far *) 1L)
            break;
    }
}
else // DOS 4, 5, 6
{
    DSKBUF4 far *dskbuf4 = *((DSKBUF4 far * far *) (doslist + 0x12));
    WORD seg, first, in_use;

    if (_osmajor >= 5) // DOS 5+ disk buffer info struct
        dskbuf4 = *((DSKBUF4 far * far *) dskbuf4),

        seg = FP_SEG(dskbuf4),
        first = FP_OFF(dskbuf4);
    in_use = 0;

// code added in case BUFFERS in HMA, have to check A20
{
    DWORD lin = (DWORD) seg << 4L;
    lin ^= first;
    if (lin > 0x100000L)
    {
        DWORD far *low = (DWORD far *) 0x0000080L;
        DWORD far *high = (DWORD far *) 0xffff0090L;
        puts("BUFFERS are HIGH");

        if (testing)
            xms_local_disable_a20(); // force A20 to test code below

        if ((low[0] == high[0]) && (low[1] == high[1]))
        {
            puts("A20 is off, turning A20 on");

            if (! xms_local_enable_a20())
                fail("Can't enable A20");
        }
        else
            // as we know from chapter 6, this will work too!
            // in fact, we could just do this, and let DOS do all checking
            _asm mov ax, 5000h
            _asm int 21h

        if ((low[0] == high[0]) && (low[1] == high[1])) // still?
            fail("A20 still disabled!");
    }
}

}

for (buffers=1, ; buffers++)
{
    if (dskbuf4->drive == 0xff)
        printf(" %p -- unused\n", dskbuf4),
    else
    {
        in_use++,
        printf("%p -- %c: %llu -- %s\n",
            dskbuf4,
            'A' + dskbuf4->drive, // pointer
            dskbuf4->sector, // drive
            buff_status(dskbuf4->flags)); // sector number
        // status
    }
    if (dskbuf4->next == first) // LRU list wraps around

```

```

        break,
    else
        dskbuf4 = (DSKBUF4 *ar *) MK_FP(seg, dskbuf4 + next),
    }
    if (buffers != doslist[0x3F]) // number of buffers
        fail("Something wrong"),
    printf("%2u buffers in use\n", in_use);
    return buffers,
}
}
}

main(int argc, char *argv[])
{
    WORD num_buf;
    if (argc > 1) testing++; // any command line arg
    // To make program useful, stick various experiments in here
    // (e.g., call 21/00 to flush disk, read/write a file, etc.)
    num_buf = buffers();
    printf(" BUFFERS: %d\n", num_buf),
    return num_buf, // can check with ERRORLEVEL
}
}

```

For machines with sufficient memory, SmartDrive or another disk cache provides better performance than relying entirely on DOS buffers. When using SmartDrive, Microsoft rightly suggests using a value for `BUFFERS` value. Whereas the DOS buffers cache DOS logical sectors from floppy diskettes as well as hard disks, SmartDrive works at the INTELish level and only cares about hard disks. SmartDrive and other DOS file caches do not cache data from any devices that do not call down to INTELish. This includes CD-ROM drives, for example, hence the growing popularity of CD-ROM "speedup" programs. As this book went to press, Microsoft announced that SmartDrive 4.2 will cache only CD-ROM drives well as included with MS-DOS 6.2. SmartDrive 4.2 is also available separately from Microsoft's CompuServe forum.

After DOS 6.0 was introduced, many supposed DoubleSpace problems reported by *John Hold* actually turned out to be "cockpit errors" and user problems with using the newer SmartDrive, which introduced a "write behind cache" feature that writes disk writes possibly too soon. The typical problem scenario is that a COPY command appears to complete, the machine is sitting at the DOS C prompt and taking a bar as a sign, the user enters the machine off a COPY's directory writes haven't yet completed, very possible given SmartDrive's aggressive caching, the end result is typically reported as CHKDSK is not workable. Part of the problem was that users were not educated to check their hard drive right before turning off the machine. As part of consulting these problems, SmartDrive 4.2 doesn't guess that the DOS prompt and disk writes have completed.

For an excellent though somewhat dated discussion, see Geoff Cunniff's "Understanding SmartDrive: *The How, Why and When*" (January 1992), also see the brief discussion of "Undocumented SmartDrive" in Chapter 1. Some other Microsoft disk caches also support the undocumented SmartDrive interface.

## SysVars, or The List of Lists

Since the introduction of `CONFIG.SYS` with MS-DOS version 2.0, the DOS kernel has maintained a collection of pointers and variables near the start of its data segment. Since this collection of pointers is not officially documented, it's known by several names. The name in the DOS source code is `SystemVars`, the structure is built by the DOS initialization code, called `System` (see Chapter 6). This is often shortened to `SysVars`, which is the name used throughout this book. In the first edition, we referred to `SysVars` using the biblical sounding name, List of Lists.

SysVars together with INT 21h AH=52h, which returns in ES:BX a pointer to this structure is the central clearinghouse for virtually all of the undocumented data concerning the DOS file system—in addition, as already discussed in Chapter 7, SysVars provides the start of the Memory Control Block chain and the device driver chain. More than any other single structure, SysVars is the key to reaching the undocumented areas of DOS. The following is a schematic listing of just some of the structures that you can access directly or indirectly via SysVars. You can see why this is sometimes called the List of Lists!

```

Memory Control Block (MCB)
  Program Segment Prefix (PSP)
    Environment segment
    Job File Table (JFT)
Drive Parameter Block (DPB)
  File Allocation Table (FAT)
  Directory entries
System File Table (SFT)
Device driver chain
Disk buffers
Current Directory Structure (CDS)
FCB table
SHARE, EXE hooks
  
```

These structures are interconnected. For example, the SFT entry for block devices contains a pointer to the corresponding DPB, so does the CDS. One of the items contained in the DPB is a pointer to its corresponding device driver. Meanwhile, the heads of both the DPB and device chains are found directly in SysVars.

It is SysVars that binds SysVars on the fly each time you boot your system. The starting point is the processing of CONFIG.SYS, which makes any installable device drivers part of the DOS kernel and possibly modifies certain values (for example, `LASTDRIVE` and the CDS pointer) stored in SysVars.

But what happens if no CONFIG.SYS file exists? Obviously no drivers are installed, but if a default values that control the building of the SysVars are assembled into IO.SYS. Thus DOS sets `FILES=8`, `LASTDRIVE=B`, `FCBS=4`, calculates an appropriate `BUFFERS` value from the memory size and drive data, and sets the primary shell to `C:\COMMAND.COM` (assuming C: is the boot drive).

If CONFIG.SYS is processed, these commands will overwrite the default values (Incidentally, CHKDSK appears to have new commands such as `FILESHIGH` and `LASTDRIVEHIGH`.) When DOS has parsed the entire file, with all commands executed or passed over with error messages, DOS builds SysVars from the values which then exist in the CONFIG control variables (see Chapter 7). It then discards the control variables along with the rest of the raw-supply initialization code.

Coincidentally, addition of the capability to load DOS high (into the High Memory Area) and to place device drivers and TSRs into Upper Memory Blocks has had some strange side effects on what happens during IO.SYS initialization. Before DOS 5.0 it was possible to predict just where in RAM each item processed from CONFIG.SYS would wind up. If high or upper memory is involved, however, you can't predict the memory map layout at all, because it depends on the availability of UMBs and the HMA which, in turn, depend on the specific device drivers installed, such as HIMEM.SYS, ENH386.EXE, or 386MAX.SYS. That is, the addition of a single driver can now totally rearrange the locations in RAM at which all other drivers are loaded, moving some from high to low and others from low to high. For example, if CONFIG.SYS requested DOS=HIGH but an XMS server such as HIMEM.SYS was unavailable, you could end up with a rather strange memory map.

### The Current Directory Structure (CDS)

SysVars contains a far pointer to an array of Current Directory Structures. Each drive on the system has its own CDS, which contains the current working directory and points to the DPB for that drive. This structure also contains two bits that specify whether the drive exists or not, whether it is modified by the JOIN or SUBST commands, or whether it is a network drive. Microsoft adopted the CDS as part of the networking additions begun in DOS 3.0; it plays a central role in manipulating foreign (not just network) file systems.

The CDS is a compact structure CDS for each possible disk device or drive letter on the system. If you specify LANSERVER as a CURDIR value, your system will have a 26-element CDS array, using the default value for LANSERVER as the CDS array's common only five elements. In other words, the LANSERVER value is nothing more than the size of the CDS array. Each element of the array is 8154 bytes long under DOS version 3.0, 8388 bytes for versions 4.0 and up. CURDIR.H Listing 8-9 illustrates a CDS structure and several delete programs later in this chapter will use. As you can see, the CDS for each drive starts off with a 67-byte ASCII string for the current path on the drive; it is from this that the CDS takes its name.

#### Listing 8-9: CURDIR.H

```
/*
CURDIR.H
Current Directory Structure (CDS), also Stacker and DoubleSpace stuff
Andrew Schuman, revised July 1993
*/

typedef unsigned char BYTE;
typedef unsigned short WORD;
typedef unsigned long DWORD;
typedef int BOOL;

#include "diskstuf.h" // for get_dpb(), DPB structure

#pragma pack(1)
typedef struct {
    BYTE current_path[67], // current path 00h
        WORD flags, // NETWORK, PHYSICAL, JOIN, SUBST, CDROM 43h
        DPB far *dpb, // pointer to Drive Parameter Block 45h
        union {
            struct {
                WORD start_cluster, // root 0000, never accessed: FFFFh 49h
                DWORD unknown,
            } LOCAL, // if ( (cds[drive] flags & NETWORK)
            struct {
                DWORD redirfs_record_ptr,
                WORD parameter,
            } NET, // if (cgs[drive].flags & NETWORK)
        } u,
        WORD backs_ash_offset, // offset in current path of '\ ' 4fh
        // DOS4 fields for IFS
        // 7 extra bytes...
    } CDS;

// flags (CDS offset 43h)
#define NETWORK (1 << 15)
#define PHYSICAL (1 << 14)
#define JOIN (1 << 13)
#define SUBST (1 << 12)
#define REDIR_NOT_NET (1 << 7) // CDROM

CDS far *curdir(unsigned drive),
extern void fail(const char *, // app must define
/* Stacker see IS_STACK.C */
typedef struct {
```





```

#include "currdir.h
typedef enum { UNKNOWN=-1, FALSE=0, TRUE=1 } OK;
/* return pointer to CDS for a given drive */
CDS far *currdir(unsigned drive)
{
    /* statics to preserve state; do one-time init */
    static BYTE far *dir = (BYTE far *) 0;
    static OK ok = UNKNOWN,
    static unsigned currdir_size,
    static BYTE lastdrv;

    if (ok == UNKNOWN) /* do one-time init */
    {
        unsigned drv_ofs, lastdrv_ofs,
        /* curr dir struct not available in DOS 1.x or 2.x */
        if (! (ok = (_osmajor >= 3)))
            return (CDS far *) 0,

        /* compute offset of curr dir struct and LASTDRIVE in DOS
        list of lists, depending on DOS version */
        #define DOS(maj, min) ((_osmajor == (maj)) && (_osminor == (min)))
        if (DOS(3,0)) { drv_ofs = 0x17; lastdrv_ofs = 0x18; }
        else { drv_ofs = 0x16; lastdrv_ofs = 0x21; }

        _asm    push si /* must preserve */
        /* get DOS list of lists into ES:BX */
        _asm    mov  ah, 52h
        _asm    int  21h

        /* get LASTDRIVE byte */
        _asm    mov  si, lastdrv_ofs
        _asm    mov  ah, byte ptr es [bx+si]
        _asm    mov  lastdrv, ah

        /* get current directory structure */
        _asm    mov  si, drv_ofs
        _asm    les  bx, es [bx+si]
        _asm    mov  word ptr dir+2, es
        _asm    mov  word ptr dir, bx
        _asm    pop  si

        /* OS/2 DOS box sets dir to FFFF:FFFF */
        if (dir == (BYTE far *) -1L) ok = FALSE,

        /* compute curr directory structure size */
        currdir_size = (_osmajor >= 4) ? 0x5B : 0x51,

    }

    if ((ok == FALSE) || (drive >= lastdrv))
        return (CDS far *) 0;
    else /* return array entry corresponding to drive */
        return (CDS far *) &dir[drive * currdir_size];
}

```

Like most of the `LASTDRIVE` programs in Chapter 2, `currdir` uses offsets computed at run time rather than C data structures set at compile time, because this seems better suited to the volatility of undocumented DOS. The code assumes that DOS 5.0 and higher are fairly compatible with DOS 4.0. There are problems with this assumption, though, since while the DOS box in OS/2 1.10 presents itself as a program as DOS 10.10, in fact it more closely resembles DOS 3.x than DOS 4.x. The test for `!_osmajor + _osminor >= 4` groups the OS/2 DOS box together with DOS 4.x instead of DOS 3.x. However, the DOS box in OS/2 1.x and 2.x don't provide a CDS anyway, and `currdir` anticipates this possibility by checking for the invalid `!!` pointer `FFFF:FFFF`.



Offset 45h from the start of the CDS entry holds a far pointer to the device DFB `cds[drive]dph`. Since the DFB in turn contains a far pointer to the actual device driver, this links the logical unit to the physical drive. Earlier in this chapter, we noted that `get_dph_drive` hits the disk, where `cds[drive]dph` appears as 0. On the other hand, for removable media, which is where hitting the disk is a major problem in the first place, `cds[drive]dph` may hold stale data.

For local drives that is where `cds[drive]flags & NETWORK` (offset 49h) in the CDS contains the starting cluster for the current directory, `cds[drive]DCM` starts cluster DCS of course. For network ones, it is 0; other programs could use this for a different purpose, to implement the carrier NAME.U.S. program, say, a stage 8 file for directories, with an overlay to traverse the directory structure and M.S. (a program could CD to a directory and pull the cluster out of CDS, citing DCS, do the hard work. You will see later that the SEI similarly contains the starting cluster for open files).

The first entry in `offset 0` of each device CDS is the current path, `cds[drive]current_path`, but it is not always what you might expect. The current path in the CDS is by convention where the data really is, rather than a device address. You can see this in the output from a `status` saved `ENUMDRV` (the source code for which you will see shortly).

```
C:\NUNDOC2>\dos\join at c:\floppy
C:\NUNDOC2>subst @: c:\nundo2
C:\NUNDOC2>rem e: and f: are network drives
C:\NUNDOC2>phantom -5512 h.
C:\NUNDOC2>enumdrv
A C:\FLOPPY          JOIN [ROOT]
B 0:\               EMOT_ACC
C C:\NUNDOC2       [F1B] STACKER (swapped D.)
D 0:\NUNDOC2       [2654B]
E \\BIBL\EXPORT\DOS NETWORK
F \\MORE\U\ANDREW NETWORK
G C:\NUNDOC2       SUBST [EMOT_ACC] STACKER
H Phantom H:\      NETWORK
```

The `COMMAND.COM` prompt shows we were logged into drive C: at the `1 NUNDOC2` sub-directory. This shows up correctly in the CDS for drive C. The `SUBST` command permits us to address files on the drive as if we were in another volume, as this `rem e: and f: are network drives`. The situation is reverse (when you use `JOIN` to refer to an entire drive as though it were a subdirectory of another drive, i.e. the CDS for drive A, always). The first byte of the current path string contains the drive letter of the `SUBST` or `JOIN` target.

As you can see, drives that are in the `ENUMDRV` map above, network drives are treated differently from local drives. Attaching a drive to a network file server means that all references to its drive actually refer to the server, which is generally thought of as always addressed with its opening "N" string. It is a special program. The prefix is part of the so-called Universal Naming Convention (UNC) that Microsoft uses, is different in networking projects such as PC File Manager and Windows 3.1 Workgroups. Here, my net PC (PC from ELP Software, drives L and I were mapped to different servers) on a Sun SPARC station running SunOS. This is a good illustration that it is the CDS files, not a copy installed by the systems. You can use the disk of a UNIX RIM machine as though it were part of DOS.

The word `offset 0` of the CDS `cds[drive]backslash_offset` contains the number of characters in the path string that precedes the root directory indicator. This is often 0 (it is set to a value of 2 to skip the escape character, i.e. when a `SUBST` command is processed, the various `g`es to skip over the `cds[drive]current_path` directory names come back by `SUBST`). That is, `Subst C:\CDOS` implies the same link DCS into `cds[drive]current_path` as `cds[drive]DCM` (the starting cluster pointer) from drive C, and sets the `SUBST` bit in the drive C's status word to set the directory

cluster number to that for the first sector of the CDFS directory on drive C, and sets the word at 4fh to a value of 1—the number of characters preceding the final \ of the pathspec string.

The CDFS may store the non-DOS name of an alien file system, such as JINX\FATFOR\INDOS or HOME\CAN\PRJ\W or the ENUMDRV output above. Backslash offset can also block off such names. The Phantom sample program later in this chapter uses this (see the ENUMDRV output). But this is possibly not normal behavior for a network redirector. An engineer at Novell, whose JRCDFS does not quite support this behavior, has told us that no other network redirector uses the CDFS to negate this way. Perhaps Phantom should instead use UNC naming.

As you navigate the directory tree, the path string stored in the CDFS tracks your position so that JRCDFS can also convert your relative path references (those which do not begin with a backslash) into fully qualified path names. Whenever you change directories by calling INT 21h function 3Bh or its user equivalent, the D command, JRCDFS updates cbs\_drive, current\_path.

Conversely, changing cbs\_drive, current\_path instantly changes the current directory. Going into your favorite debugger, locating the CDFS for the current drive, and manually editing the path string in the CDFS is sufficient to change directories. The change is immediately reflected in the PROMPT SpSp display, for example. It's now clear why this is called the Current Directory Structure.

However, this is one problem with the CDFS array. There's only one. All DOS tasks share the single cbs\_CDS. The reader who thinks of DOS as a single-tasking operating system might well ask, "How DOS tasks in DOS each do only one task, so of course there is only one CDS?" But as shown in Chapter 7's discussion of PMN and Chapter 9's discussion of VSRs and DOS multitasking, this isn't so. Consider especially the example of Windows Enhanced mode, where multiple DOS boxes can be, practically, multitasked, each manipulating the DOS current directory at the same time. Windows uses instance data to maintain separate CDS for each DOS box (see Chapter 1). Further, now, each Windows program has its own current drive and directory, maintained in its Task Database (see *Undocumented Windows*, Chapter 5).

### Walking the CDS Array

The article mentioned in Listing 8.10 goes to some trouble to ensure that it can be called frequently without any duplicated effort. This way, curdir, can be called as a loop for each drive in the system, producing the ENUMDRV output shown earlier. Listing 8.11 shows ENUMDRV.C, which contains the actual loop along with some other code we'll get to in a moment.

#### Listing 8.11: ENUMDRV.C

```
/*
ENUMDRV.C -- uses curdir() in CURDIR.C
Added CD ROM bit, DoubleSpace test, Stacker test, primitive RAM disk test
Andrew Schulman, revised July 1993
*/
bcc enumdrv.c curdir.c is_stack.c is_dupac.c diskstuff.c
*/

#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <dos.h>

#include curdir.h // includes DoubleSpace, Stacker too
#include "diskstuff.h" // get_dpb

void fail(const char *s) { puts(s); exit(1); }

int maybe_ram_disk(int drive);

main()
{
    CBS far *dir;
    BYTE host;
    BOOL swapped,
```

```

unsigned lastdrv;
int i, seq;
dos_setdrive(DzFF, &lastdrv);
for (i=0, i<lastdrv, i++)
{
    if (! (dir = currdir(i)))
        fail("can't get current directory structure");
    else if (dir->flags) /* is this a valid drive? */
    {
        printf("%c 2-4Dfs ", 'A' + i, dir->current_path);

        if (dir->flags & JOIN) printf("JOIN ");
        if (dir->flags & SUBST) printf("SUBST ");
        if (dir->flags & REDIR_NOT_NET) printf("REDIR_NOT_NET ");
        if (dir->flags & NETWORK) printf("NETWORK ");
        else switch (dir->u LOCAL.start_cluster)
        {
            case 0: printf("[ROOT] "); break; // root dir
            case DzFFFF: printf("[NOT_ACC] ", break; // not accessed
            default: printf("[%u] ", dir->u LOCAL.start_cluster);
        } // check cluster numbers by running CLUSTMAN
        if (!double_space_drive(i, &swapped, &host, &seq))
            printf("DBLSPACE %c:\\dblspc (%Su", 'A' + host, seq);

        if (stacker_drive(i, &swapped, &host))
        {
            printf("STACKER");
            if (swapped) printf(" (swapped %c)", 'A' + host);
        }

        if (maybe_ram_disk(i)) printf("RAM disk");

        putchar('\n');
    }
}
return 0;
}

int maybe_ram_disk(int drive)
{
#ifdef USE_DPS_BPB
    return (currdir(drive)->dpb->copies_fat == 1);
#else
    DPS dpb = get_dpb(drive+1);
    return (dpb) ? (dpb->copies_fat == 1) : 0;
#endif
}

```

For each drive FASTDRIVE\_ENUMDRV calls the currdir function and prints out values from the CDS. For s the current path string the drive type and for lsd drives the starting cluster of the current directory. ENUMDRV then calls some additional functions to check for DoubleSpace and Stacker compressed drives and for RAM disks. We'll get to the DoubleSpace and Stacker code in a moment.

### Detecting RAM Disks

The maybe\_ram\_disk function is interesting. Since RAM disks are intended to look as much like physical disks as possible, there's no sure test for a RAM disk. However, RAM disks generally provide only one FAT, and the presence of a single FAT is a common test for a truly RAM disk. As noted earlier, for floppy and hard disks DOS forces the DPB to indicate two FATs, even when the BPB says

The number of FATs is retrieved from the DPR (note that network drives don't have DPRs). The major disk function can get the DPR either from the `C:\DOS\currentdrive >dpr` or by calling the `get_dpr` routine from `DISKINFO.C` (Listing 8-4). This involves the usual trade-off between letting the disk write `get_dpr` and possibly getting a stale DPR with `currentdrive >dpr`. Listing the disk write `get_dpr` isn't so bad because this function checks for removable media and contains a critical error handler.

For an alternative approach to drive enumeration, see the `DRVINFC` program in Chappell's *DOS Internals*.

### DoubleSpace Drives

DoubleSpace is an on-the-fly disk compression subsystem that Microsoft licensed from Veritas, makers of the earlier DoubleDisk product, and incorporated in MS-DOS 6.0. `ENUMDRV` checks for DoubleSpace drives by calling the `double_space_drive` function from `IS_DSPACE.C`, shown in Listing 8-12.

### Listing 8-12. IS\_DSPACE.C

```

/*
IS_DSPACE.C -- Is this a DoubleSpace drive?
Based on DRVINFC from Microsoft's DSDUMP sample code (March 1993)
Andrew Schuman, July 1993
NOTE: *None* of this is undocumented!
*/
#include <stdio.h>
#include <stdio.h>
#include <dos.h>
#include <cardr.h>

BOOL DSGetDriveMapping(int drive, BOOL *pcomp, int *phost, WORD *pseq)
{
    BYTE ldata, hdata;
    _asm mov ax, 4A11h /* DBLSPACE INT 2Fh function */
    _asm mov bx, 1 /* DSGetDriveMapping */
    _asm mov di, byte ptr drive
    _asm int 2Fh
    _asm r ax, ax
    _asm jnz no_dblspace
    _asm mov byte ptr ldata, bl
    _asm mov byte ptr hdata, bh
    *pcomp = ldata & 0x80, /* compressed drive flag */
    *phost = ldata & 0x7F, /* host drive (if drive compressed) */
    *pseq = hdata, /* compressed drv sequence # (0..254) */
    return 1;
no_dblspace:
    return 0;
}

// Can generate CVF filename with host \DBLSPACE seq
BOOL double_space_drive(BYTE drive, BOOL *pswap, BYTE *phost, int *pseq)
{
    WORD seq, seq2,
    int drHost, drHost2,
    BOOL fCompressed, fSwapped;
    if (! DSGetDriveMapping(drive, &fCompressed, &drHost, &seq))
        return 0;
    if (! fCompressed)
        return 0; /* could be host drive, but don't care */
    fSwapped = 0;
    if (DSGetDriveMapping(drHost, &fCompressed, &drHost2, &seq2))
        if (drHost2 == drive) // host of host is drive itself: means swapped
            fSwapped = 1;
}

```



How do some utilities need to know about the actual DoubleSpace structures such as the MDA1 and MDA2? For example, Norton Utilities 7.0 includes versions of the Norton Disk Doctor, NDD, and Speedy that can repair and defragment DoubleSpace and Tracker drives, presumably in a format that Microsoft's DISKPART, CHKDSK, or Star Electronics' SC-HFCK utilities can't.

As an application of DoubleSpace information, consider writing a program to compute the compression ratio. For example, DIR /C does in DOS 6.0:

```
Volume in drive A is DOS6
Volume Serial Number is 1603-265C
Directory of A:\

COMMAND.COM      52925 03-10-93   6 00a  1 4 to 1 0
README.TXT       350 07 12 93   3 11p 16,0 to 1,0
CHAPS.ZIP        103029 07 16-93   6,15p  1 0 to 1,0
CHAPSC.ZIP       58469 07-16-93   2 12p  1 3 to 1,0
DR.ZIP           24075 07-10-93  10 27a  1 0 to 1 0

****
FAT             8270 07-10-93   9 26p  2,5 to 1 0
DPBTEST.EXE    8354 07-15-93  12,18p  2 5 to 1 0
              1.7 to 1.0 average compression ratio
49 File(s)      550393 bytes
                1384448 bytes free
```

How does DPBTEST know that CHAPS.ZIP has 1:1 compression, while DPBTEST.EXE has 1.7:1 compression? Certainly the information doesn't come from any of the standard DOS file system structures. The answer comes from the MDA1. Which is not compressed DOS disks have a fixed for each cluster a cluster number to sector numbers, see CLUSTER TO SECTOR in Listing 8.5. DoubleSpace does have a cluster to sector numbers per cluster, depending on the possible compression. So, a cluster can have a cluster to sectors formula. DoubleSpace uses the MDA1, which is a table of cluster to cluster number. MDA1 is a very poor name for the structure, given that it is a table of cluster to sector numbers. MDA1 is a very poor name for the structure, given that it is a table of cluster to sector numbers. For each cluster, the MDA1 contains a four-byte entry showing the following C structure with but a few adjustments to Microsoft's C VEH header file. Note that Borland C++ does not accept the use of a C struct in practice, so you should access the MDA1 fields with shifts and masks, as in Listing 8.13.

```
typedef struct {
    unsigned long sectorStart; // starting sector for cluster
    unsigned reserved; // 0
    unsigned cComp; // 4; // # of compressed sectors
    unsigned orig; // 4; // # of uncompressed (original) sectors
    unsigned flags; // 2; // Is entry in use? data compressed?
} MDFAT_ENTRY;

MDFAT_ENTRY MDFAT[num_clusters];
```

Because 4 bits each are devoted to the count of original sectors,  $cl_{orig}$ , and the count of possibly compressed sectors found in the  $CMP_{cComp}$ , there are a maximum of 16 sectors per cluster. Given a cluster whose original (uncompressed) sectors is 16 sectors,  $cl_{orig} == 16$ , the compressed cluster occupies less than 16 sectors. With 2:1 compression,  $cComp == 8$ . With no compression (for example, a portion of CHAPS.ZIP),  $cComp == 16$ .

There is a limit on the number of clusters. With a maximum of 16 sectors per cluster and 512 bytes per sector, the DOS space can only hold a maximum of 512 megabytes of *uncompressed* data.

The  $cl_{orig}$  field of the MDA1 entry is important. DoubleSpace does not assume a fixed 16 uncompressed sectors per cluster. It merely assumes the maximum. Files whose true size is not a multiple of 16 sectors will have a few MDA1 entries whose  $cl_{orig}$  field is less than 16. In this way, DoubleSpace changes the space usage of cluster overhang of hard disks under DOS to a much less wasteful sector overhang.



You can see this clearly in the following session using two of our earlier programs, NAMELIST (see Listing 8-5) and FATLSEX (Listing 8-2), as well as a Microsoft DoubleSpace dumper called `appropriate enough DSDUMP.MP`. Microsoft provides the C source code for DSDUMP along with the DoubleSpace documentation on its CompuServe forum:

```
C:\MSDOCS2>namelist a \command.com
A:\COMMAND.COM ==> 12 (sect 256, 0x000000ec)
C:\MSDOCS2>fat a: 12
12-18 (7)
7 clusters in 1 groups
C:\MSDOCS2>dtdump a: /M14-20
DoubleSpace File Dumper - Version 0.58
Drive: A (mounted from I:\DBLSPACE\000)
MDFAT entries 14 to 20
Flags: A=Allocated, F=free, C=Compressed, U=Uncompressed
-----
FAT#  MDFAT#  Flags  cunc  cComp  secStart
-----
 12     14     A,C    16    13     261
 13     15     A,C    16    14     274
 14     16     A,C    16    15     288
 15     17     A,C    16    14     303
 16     18     A,C    16    12     317
 17     19     A,C    16    10     329
 18     20     A,C     8     1     339
```

This shows where DIR-C gets those compression ratios. For `COMMAND.COM`, recall that DIR-C shows a compression of 14 to 10. Adding up the `cunc` column, we get  $16*6=96$  sectors for the original `COMMAND.COM`. Adding up the `cComp` column, we get  $13+14+15+14+12+10+1=79$  sectors used. If `DBLSPACE=000`, dividing 104 by 79, we get 1.31.

But if the actual compression ratio is 1.41, why does DIR-C get 1.4? Because, rather than add up the actual cluster sectors, it uses the `File` (taking into account the total count that is generally less than 16) for `FAT #18` (it was 8). DIR-C instead simply multiplies the number of clusters by 16:

$$\text{compression ratio} = (\text{num\_clusters} * 16) / \text{total\_cComp}$$

This formula explains why DIR-C shows 16.0 to 1.0 compression ratios for all files whose size is less than 512 bytes. Stack's SDIR shows similarly wildly inflated compression ratios for small files. With DIR-C, even a file whose size shows a compression ratio of 16.0 to 0.0 (it's been squeezed down to 16 bytes), this winds completely bogus at first, but in fact it is only mostly bogus. As we pointed out earlier, a file whose size is less than 512 bytes still occupies the full 16-sector on an uncompressed disk. By multiplying the number of clusters by the sectors per cluster, DIR-C reflects the fact that DoubleSpace compresses 16 sectors of data into one sector, but the original file occupies a total of 16 sectors of space on the original storage device (i.e., a single sector).

The only problem is that DIR-C assumes 16 sectors per cluster, which is appropriate for the DoubleSpace files, but probably not appropriate for the original uncompressed media. As pointed out earlier, by default, 5.25-inch floppy disks use 1 sector per cluster, and hard disks generally use 4 or 8 sectors per cluster. Thus, a file whose true compressor ratio is 1.1 for a floppy disk or 4.1 or 8.1 for a hard disk will still look like 16.1 for both media with 16 sectors per cluster. More accurate compression ratios require adding a file's size (cluster size) to the sectors per cluster from the DIR-C file, original or DoubleSpaced form:

$$((\text{num\_c\_clust} - 1) * 16) + \text{original\_sect\_per\_clust}) / \text{total\_cComp}$$

DIR-C uses a similar formula, producing more accurate compression ratios.

		bytes	/C ratio	/CM ratio
COMMAND.COM	COM	52925	1.4	1.3
ENUMDRV	EXE	9252	2.1	1.3
READTHIS	TXT	350	16.0	1.0
DN	ZIP	24075	1.0	-1.0
DSDUMP	ZIP	13956	1.0	1.0
CHAPB	ZIP	103029	1.0	1.0
CHAPBC	ZIP	58469	1.1	1.0
FILES	C	7949	1.8	1.8
DPBTEST	C	1957	5.3	1.3
FAT	EXE	8270	2.5	1.5
DPBTEST	EXE	8354	2.5	1.5
Average compression ratio			1.4	1.1

For example, COMMAND.COM was located on a high density floppy with one sector per cluster, so that a 4:1 compression ratio was accurate. The true compression ratio—that is, the effect of CMT on DoubleSpace vs. not using it—for a file like READTHIS.TXT is not 16:1 but instead depends on the original pre-DoubleSpace media. For hard disks it is 8:1 or 4:1 for high density floppy and is only 1:1.

If we're therefore still a problem because—at least in the original release of DoubleSpace, the number of compressed sectors can in some cases exceed the original number of uncompressed sectors on the host media (that is,  $c \times mp > c \times mp$ ). For example:

```
C:\UNDOC2>namclust f:dsdump.zip
A:\DSDUMP.ZIP ==> 5 (sect 36, D=00000024)
C:\UNDOC2>fat f: 5
5-32 (28)
28 clusters in 1 groups
C:\UNDOC2>namclust a:dsdump.zip
A:\DSDUMP.ZIP ==> 32 (sect 556, D=00000022c)
C:\UNDOC2>fat a: 32
32-33 (2)
2 clusters in 1 groups
C:\UNDOC2>dsdump a: /R34-35
DoubleSpace File Bumper - Version D.58
Drive A (mounted from I:\DBLSPACE 000)
MFAT entries 34 to 35
Flags: A=Allocated, F=Free, C=Compressed, U=Uncompressed
-----
FAT#  MBFAT#  Flags  cUmc  cCmp  secStart
-----
    32     34  A,U  16  16      454
    33     35  A,U  12  16      470
```

In this case, the original occupied 28 one-sector clusters and the “compressed” version occupies only 16 sector clusters, so the compression ratio is 28 to 16, or 0.875 to 1.0. Both DIR /C and DIR /CH report a 1:1 compression ratio. The problem here isn't a faulty compression algorithm. As the /F flag above shows, DoubleSpace has decided it can't compress this data, so it is stored uncompressed. It is not clear why DoubleSpace should bother storing 16 sectors when the original fits on 12. One developer tells us that this is just a bug that is being fixed, and that we're making too big a deal of it here.

In any case, Microsoft's DSDUMPMP shows that you can derive compression ratios (whether true compression ratios or the inflated ones DIR /C produces) by examining the MFAT entries in the CMT.

How then does a program find the MIDEA1 for a DoubleSpace drive? First call the DoubleSpace DSKCsdvcsMapping API function INI 2Fh AX 4411h to find the host DBLSPACE1 vsq CVI file name, as shown by the DoubleSpace driver function IS\_DSPACE (Listing 8-13). Next open the CVI file and read in the MDRBPB to find the MIDEA1; the MDRBPB is an extended BBP at the start of the CVI file and includes a secMIDEA1Start field.

The sample program in Listing 8-13, when run with a VERBOSI command-line switch, outputs the file's MIDEA1 entries from a CVI file file.cvi itself; this program is not terribly useful, as you could list it to the `dir` command, NAMLIST.FAT, and IS\_DSPACE to produce a program to show accurate compression ratios for any DoubleSpaced file. Use the code from NAMLIST to turn a filename into a starting cluster, use the FAT code to test the file's entire set of clusters, use IS\_DSPACE to locate the CVI file on the host drive, and then use the code in Listing 8-13 to examine the MIDEA1 entries for the file's clusters.

### Listing 8-13. MDFAT.C

```

/*
MDFAT.C -- Dump a DoubleSpace MDFAT
Based on Microsoft's DSDUMP sample code
Andrew Schulman, July 1993
NOTE: *NONE* of this is undocumented!
cl mdfat c is_dspac c
*/

#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <ctype.h>
#include <dos.h>

typedef unsigned char BYTE;
typedef unsigned short WORD;
typedef unsigned long DWORD;
typedef int BOOL;

#pragma pack(1)
typedef struct {
    BYTE    jmpBOOT[5]; // Jump to bootstrap routine
    char    achOEMName[8]; // OEM Name ("MSDOS6 0")
    // MS DOS BPB
    WORD    cbPerSec;
    BYTE    csecPerClu;
    WORD    csecReserved;
    BYTE    cFATS;
    WORD    cRootDirEntries, csecTotalWORD;
    BYTE    bMedia;
    WORD    csecFAT, csecPerTrack, cHeads;
    DWORD   csecHdden, csecTotalDWORD;
    // DoubleSpace extensions
    WORD    secMDFATStart;
    BYTE    nLog2cbPerSec;
    WORD    csecMDReserved, secRootDirStart, secHeapStart, clufirstData;
    BYTE    cpageBIFAT;
    WORD    RESERVED1;
    BYTE    nLog2csecPerClu;
    WORD    RESERVED2;
    DWORD   RESERVED3, RESERVED4;
    BYTE    f12BitFAT;
    WORD    cmbCVFRax;
} MDRBPB;

#pragma pack()

void fail(const char *s) { puts(s); exit(1); }
// my buffer size

```

```

#define MFAT_PER_BLOCK      (2048)
#define MDFAT_BLOCK_SIZE   (MDFAT_PER_BLOCK * sizeof(DWORD))
// handy macros from Microsoft's OSDUMP.C
#define GET SECSTART(dw)    ((dw) & 0x3FFFFFF)
#define GET CSEC_CODED(dw)  (1 + (int) (15 & ((dw) >> 22)))
#define GET CSEC_PLAIN(dw) (1 + (int) (15 & ((dw) >> 26)))
#define GET_FLAGS(dw)      (3 & ((dw) >> 30))
#define USED(dw)           (GET_FLAGS(dw) & 2)
#define UNCOMPRESSED(dw)   (GET_FLAGS(dw) & 1)

static DWORD sectors_coded = 0, sectors_plain = 0;
static WORD clusters_in_use = 0;
static WORD first_data_clust = 0;
static int verbose = 0, did_banner = 0;

void do_mdfat_entry(WORD cluster, DWORD mdfat_entry)
{
    int coded, plain;
    if (USED(mdfat_entry)) clusters_in_use++;
    else return;

    sectors_coded += (coded = GET_CSEC_CODED(mdfat_entry));
    sectors_plain += (plain = GET_CSEC_PLAIN(mdfat_entry));

    if (verbose || (coded > plain))
    {
        if (! did_banner)
        {
            if (coded > plain)
                printf("Expanded clusters (cCmp > cUnc):\n");
            printf("Cluster   Sector cCmp cUnc C/U\n");
            printf("                \n");
            did_banner++;
        }
        printf("%5u %5u %5u %5u %5u %5u\n",
            cluster - first_data_clust,
            GET_SECSTART(mdfat_entry),
            coded, plain,
            UNCOMPRESSED(mdfat_entry) ? 'U' : 'C');
    }
}

main(int argc, char *argv[])
{
    char *cvf_name,
    FILE *cvf;
    MDDBPB *mddbpb;
    DWORD *mdfat_block, num_sect, mdfat bytes,
    WORD num_clust, mdfat blocks, ratio, '\n';

    if (argc < 2 || argv[1][1] == '?')
        fail_usage_mdfat ["DoubleSpace drive or CVF file"];
    if (strcmp(strupr(argv[1]), "--VERBOSE") == 0)
        { verbose++; argv++; argc--; }
    cvf_name = argv[1];

    if ((cvf = fopen(cvf_name, "rb")) == NULL)
    {
        // If we can't open the specified filename, maybe it's not
        // a filename at all. See if maybe it's a DoubleSpace drive
        // letter. If so, use function from IS_DSPAC.C to get CVF name
        extern BOOL double_space_drive(BYTE drive, BOOL *pswap,
            BYTE *phost, int *pseq);
        BYTE drive = (BYTE) (toupper(cvf_name[0]) - 'A'),
        BYTE swap, host;
        int seq;
    }
}

```

```

if (double_space_drive(drive, &swap, &host, &seq))
{
    static char filename[16];
    sprintf(filename, "%c \\dblspc.LDSu", 'A' + host, seq);
    printf("CVF file: %s\n", filename);
    cvf_name = filename;
    if (!(cvf = fopen(cvf_name, "rb"))) == NULL)
        fail("can't open CVF file");
}
else
    fail("can't open CVF file");
}
if (! (mdbpb = malloc(sizeof(MDBPB))) || fail("insufficient memory"),
    if (! fread(mdbpb, sizeof(MDBPB), 1, cvf)) || fail("can't read MDBPB"),
    if (! (mdbpb->jmpBOOT[0] == 0xE9) || (mdbpb->jmpBOOT[0] == 0xEB))
        fail("not a valid CVF file"), // has to start with JMP

first_data_clust = mdbpb->clust1rstData, // used by do_mdfat_entry()
num_sect = (mdbpb->csecTotal/WORD) *
    mdbpb->csecTotal/WORD - mdbpb->csecTotal/BWORD,
num_clust = 1 + (num_sect / mdbpb->csecPerClu),
mdfat_bytes = (BWORD) num_clust * sizeof(BWORD),
mdfat_blocks = 1 + (mdfat_bytes / MDFAT_BLOCK_SIZE),

if (! (mdfat_block = (BWORD *) malloc(MDFAT_BLOCK_SIZE)))
    fail("insufficient memory");

fseek(cvf, (1 + mdbpb->secMDFATStart) * mdbpb->cbPerSec, SEEK_SET),
for (i=0; i<mdfat_blocks-1; i++)
    if (!fread(mdfat_block, MDFAT_PER_BLOCK, sizeof(BWORD), cvf))
        for (j=0; j<MDFAT_PER_BLOCK; j++)
            do_mdfat_entry((1+MDFAT_PER_BLOCK)+j, mdfat_block[j]),
        else
            fail("can't read MDFAT");

if (!fread(mdfat_block, MDFAT_PER_BLOCK, sizeof(BWORD), cvf) // last one
    for (j=0; j<(num_clust / MDFAT_PER_BLOCK); j++)
        do_mdfat_entry((mdfat_blocks-1) + j, mdfat_block[j]),
else
    fail("can't read MDFAT"), // okay to read less

fclose(cvf),
printf("\nTotal clusters          %u\n", num_clust),

// print global counters incremented by do_mdfat_entry()
printf("In use:                    %u (XluX used)\n",
    clusters_in_use, (100L * clusters_in_use) / num_clust),
printf("Plain sectors (lunc)       %lu\n", sectors_plain),
printf("Compressed sectors (lcmp)   %lu (%u% compressed)\n",
    sectors_coded, 100L * sectors_coded) / sectors_plain),
ratio = (10L * sectors_plain) / sectors_coded,
printf("Sector compression ratio: %u.%u to 1\n",
    ratio / 10, ratio % 10);
if (sectors_coded != 0)
{
    ratio = (100L * clusters_in_use) / sectors_coded,
    printf("DIR /C compression ratio: %u.%u to 1\n",
        ratio / 10, ratio % 10);
}

// Exercise for the reader (I suddenly got very lazy here)
// To produce DIR /C use get_dob(host)->sectors_per_cluster

return 0;

```

MDFAT shows compression ratios for a CVF file and also lists any clusters where cCmp is greater than cUnc.

```
C:\UND0C2\CHAP8>mdfat acidb1space.000
Expanded clusters (cCmp > cUnc):
Cluster  Sector  cCmp  cUnc  C/U
-----  -----
    33      470    16   12   U
    46      678    16   10   U
   105     2107    16   10   U
   112     2125    16    8   U

Total clusters:          297
In use:                  179 (60% used)
Plain sectors (cUnc):    2288
Compressed sectors (cCmp) 2000 (87% compressed)
Sector compression ratio: 1.3 to 1
DIR /C compression ratio: 1.4 to 1
```

The sector compression ratio is less than or equal to DIR /C when the original media had 1 sector per cluster. To produce a DIR /C and the MDFAT program would need to keep count of partial clusters. If it could, the program would also need to use get\_dpb(hdnt) > sectors\_per\_cluster.

As with other DoubleSpace files, it's important to realize that the actual compression and decompression is only a part of DoubleSpace. It is done by something called a Microsoft Real-Time Compression Interface (MRCI) server that happens to be found inside DEFSPLIB.SYS. There can be other MRCI processes running on hosts besides DoubleSpace and other MRCI servers. MRCI includes services for compressing and decompressing arbitrary blocks of data in memory. It includes support for run-time file compression. This interface uses INT 21h AX=4470h and the BIOS service INT 1Ah AX=1000h to communicate with bare-metal MRCI servers, and is documented in Microsoft's MS-DOS 6.0 *Programmer's Reference*. Developers can license MRCI libraries for DOS and Windows free of charge from Microsoft, though Microsoft's cover letter to the MRCI license agreement warns that MRCI might be subject to a patent infringement suit from Stac Electronics.

As with other tools, to see what the status is (the company that licensed [read: gave away] the data compression technology to Microsoft for inclusion in MS-DOS 6 DoubleSpace™ sells a nice add-on product to DoubleSpace called Space Manager). According to the packaging, "We didn't share all our space secrets with Microsoft. We held back one of our crown jewels that'll make your life a lot easier." These include SuperCompress, SelectCompress, Defragment, and other nifty additions to DoubleSpace.

## Stacker Drives

Back in Listing 8-11, INT 13h DRV also checks for Stacker drives, using the function stacker\_driver from IS\_STACK.C, shown in Listing 8-14.

### Listing 8-14: IS\_STACK.C

```
/*
IS_STACK.C Detect Stacker driver, and Stacker drives
Andrew Schulman, July 1993

These calls were documented in the back of the Stacker 2.0 User's
Guide, and are also described in the Interrupt list
*/
#include <stdlib.h>
#include <stdio.h>
#include <dos.h>
#include "curdir.h"

STACKER_DRIVER far *stacker_detect(void)
{
```

```

static STACKER_DRIVER far *drv = (STACKER_DRIVER far *) 0,
char far *pbuf,
char *buf,
if (drv) return drv;
if (! (buf = (char *) malloc(1024)))
    fail("insufficient memory"),
pbuf = (char far *) buf;

// Get address of Stacker driver by making an otherwise-legal
// INT 25h call
#define STACKER_MAGIC 0x0C0C0
_asm push ds
_asm mov ax, STACKER_MAGIC
_asm lds bx, pbuf
_asm mov cx, 1
_asm xor dx, dx
_asm int 25h
_asm popf
_asm pop ds
_asm cmp ax, STACKER_MAGIC
_asm jne no_stacker
stacker
    drv = *((STACKER_DRIVER far *) &buf[4]),
    free(buf);
    return (drv->signature == 0xA55A) ? drv : (STACKER_DRIVER far *) 0;
no_stacker:
    free(buf);
    return (STACKER_DRIVER far *) 0;
}

int stacker_swapped(int drive)
{
    static STACKER_DRIVER far *driver = (STACKER_DRIVER far *) 0,
    if (! driver) driver = stacker_detect();
    return (driver) ? driver->drive_map[drive] : 0;
}

DWORD stacker_drive(int drive, BOOL *pswap, BYTE *phost)
{
    volatile DWORD drv = 0, // volatile so compiler doesn't just return 0
    static STACKER_DRIVER far *driver = (STACKER_DRIVER far *) 0;
    int host;
    if (! driver) driver = stacker_detect(), // one-time init
    if (! driver) return 0L;
    drive++,
    _asm mov ax, 4404h
    _asm mov cx, 4
    _asm mov bl, byte ptr drive
    _asm lea dx, drv
    _asm int 27h
    if (! drv) return 0L;
    drive--;
    host = stacker_swapped(drive);
    *phost = host;
    *pswap = (host != drive);
    return drv,
}

#ifdef TESTING
void fail(const char *s) { puts(s); exit(1); }
main()
{
    int i, lastdrive, swap, host,
    if (! stacker_detect())
        fail("Stacker not installed");
    _dos_setdrive(0xFF, &lastdrive);
}

```

```

for (i=0; i<lastdrive; i++)
  if (stacker_drive(i, &swap, &host))
  {
    printf("Zc-%1STACKER", i + 'A');
    if (swap) printf(" (swapped Zc:)", host + 'A');
    printf("\n");
  }
return 0;
}
endif

```

This code checks for the presence of Stacker using an otherwise invalid INT 25h call, it also VAXes for Stacker drives by making IOCTL calls to the Stacker driver. These portions of the Stacker API are listed in the INTRENT database that accompanies this book. Some partial documentation of the Stacker VTL is available from the Stac Electronics forum on CompuServe.

Earlier we saw that DIR.C in DOS 6.0 produces very odd compression ratios for small files. The same thing has happened for cases of Stacker's MDR utility.

```

C:\STACKER>mdir /s /d *
DIR - 3 00, c Copyright 1990-92 Stac Electronics, Carlsbad, CA
Volume in drive C is STACKER
Directory of C:\STACKER

FOO      BAR      1 07-20-93  10 09p  8.0:1
FOGBAR   BAR      5 07-20-93  10 19p  16.0:1
          2 file(s)  3137536 bytes free

Overall compression ratio of files listed = 10.7:1

```

Showing 8 L compression for a one byte file and 16 L compression for a five byte file is certainly wrong. The FAT-16 file had four sectors per cluster and thus uses 2,048 bytes each to store these small files. Another Stacker uses a single 512 byte sector each to store these small files, the actual compression ratio over the MDR is probably using the maximum 16 sectors per cluster of the Stacker drive, that is 16 x 4 sectors per cluster of the original host drive. But then why the reported 8 L compression for the 1 L file. This ratio is somewhat very strange and slightly dishonest about these compression ratios.

The Stacker MDR utility has an undocumented `-D` switch that dumps out diagnostics for some of the internal Stacker structures. The `Cmap` is presumably the cluster map, similar to DoubleSpace's poorly named `MJDAI`.

```

C:\STACKER>mdir /s /d mdir.exe
DIR - 3 00, c Copyright 1990-92 Stac Electronics, Carlsbad, CA
Volume in drive C is STACKER
Directory of C:\STACKER

$DIR      EXE      35689 06-03-93  3.05a
First cluster 1986
Fat 1989 Cmap 00B4C107 Extent 00B802ED
Fat 198A Cmap 00B6C11A Extent 00B904ED
Fat 198B Cmap 00B5C2A1 Extent 00BA02ED
Fat 198C Cmap 00B3C208 Extent 00BC04ED
Fat $FFF Cmap 00B1C364 Extent 000402ED

Sector Ratio = 1.0 933
Host Cluster Ratio = 0.960:1
Compressed Disk Cluster Ratio = 1.067:1
Adjusted Cluster Ratio = 1.067:1
Percent of original file size: 106.17%
          1 file(s)  3137536 bytes free

Overall compression ratio of files listed = 1.1:1

```



NDIR.EXE itself achieved such low compression here because Stac Electronics has already pre-compressed the executable using the excellent LZEXE program by Fabrice Bellard, whose signature "FAB" appears at the end of such compressed executables. Mitugu Kurizono's UNLZEXE can decompress these executables.

### Novell NetWare Drives

Because FNCMDRV simply walks the CDS, whose size is set by FASDRIVE, the program does not show Novell NetWare drives that are assigned to drive numbers that exceed the value of FASDRIVE. As noted in Chapter 4, the NetWare API includes INT 21h AH=F1h and AH=F2h functions that can be used to display NetWare drives. However, as also explained in Chapter 4, the NetWare workstation shell also provides its own versions of many standard INT 21h calls. A documented MSDOS 7.03 system call Get Assign List Entry (INT 21h AX=5F02), can enumerate NetWare drives. Of course, it also enumerates any non-NetWare network drives that DOS supports.

Likewise, a commercial version of the FNCMDRV program would also have to incorporate code such as that in Listing 8-15, NETDRV.C, which runs the Get Assign List Entry function in a loop and prints out the local and network names for each network-connected device, including printers as well as drives.

#### Listing 8-15: NETDRV.C

```
/*
NETDRY.C -- Display network connections, using 21/5f02
Andrew Schulman, July 1993
NOTE: 21/5f02 is "not" undocumented! Is NetWare support for it documented?
*/
#include <stdlib.h>
#include <stdio.h>
#include <dos.h>

typedef unsigned char BYTE;
typedef unsigned short WORD;

int _dos_getassign list(WORD index, char far *local, char far *net,
    BYTE *pavail, BYTE *pdevtype, WORD *puserval)
{
    BYTE avail, devtype;
    WORD userval;
    _asm push ds;
    _asm push di;
    _asm push si;
    _asm mov bx, index;
    _asm lds si, local;
    _asm les di, net;
    _asm mov ax, 5f02h;
    _asm int 21h;
    _asm pop si;
    _asm pop di;
    _asm pop ds;
    _asm jc error;
    _asm mov byte ptr avail, bh;
    _asm mov byte ptr devtype, bl;
    _asm mov word ptr userval, cx;
    *pavail = avail; *pdevtype = devtype; *puserval = userval;
    return 0; // success
error:;
    // return value in AX
}

main()
{
    char local[128], net[128];
```



This is the only list interested one of the authors enough to window into DOS to seek out the book. However, `DRVSET` does not actually work so well in its intended purpose of unloading MSCDEX from memory, so it is listed at the top of `DRVSET.C` in Listing 8-16.

Access to the CD-ROM is just for invalidating drive letters. With a real CD-ROM you can actually do it another way, and do so simply by turning on some bits in the flags word. DOS immediately recognizes the flags as valid, so that you can change to it and send it requests (all of which fail, of course).

```
C:\MSNDOS>e
Invalid drive specification
C:\MSNDOS>dir e
Invalid drive specification
C:\MSNDOS>dirset e not phys
NET PHYSICAL
C:\MSNDOS>dir e
Volume in drive E has no label
Directory of E:\

File not found
C:\MSNDOS>e
E:\>chkdsk
Cannot CHKDSK a Network drive
```

Simply by invalidating bits in the CD-ROM, we convince DOS that there is some way a valid drive. Since DIR does not show us how up there, it cannot check what value this has. However, this is the foundation for our simple drives with the network redirection. So simply by saying drive E is a network drive, again, the term "network" just means a two-way, standardible, system, so that you can use DOS file requests to talk to a drive on an INTEL 2102 or some other standard. Later on in this chapter you will see such a hardware. In the meantime, `DRVSET` is useful for experimenting with all drives, which are the foundation for installable file systems under DOS.

Since `DRVSET`, `DRVSET` is not so simple, it makes sense to package them in the same way as include `DRVSET.C` in Listing 8-16. The resulting program behaves differently if its symbolic target[0] in C contains the string, "DRVSET".

Because `DRVSET`, `DRVSET` is not so simple, it makes sense to package them in the same way as include `DRVSET.C` in Listing 8-16. The resulting program behaves differently if its symbolic target[0] in C contains the string, "DRVSET".

## Listing 8-16: `DRVSET.C`

```
/*
DRVSET.C Set attributes of drive (CDS entry) given on command line
Andrew Schulman, revised July 1993
```

```
NOTE This (in its DRVSET form) was the stupid program that got me
interested in undocumented API calls, back in 1989 when I was working
on CD/Netwo ker for David Moxey at Lotus. It was part of something
that our customers (especially Shearson) wanted to unload MSCDEX,
called NOMSCDEX. When Shearson talks, Lotus listens. Basically,
NOMSCDEX was a batch file that ran MAKE/RELEASE and then DRVSET.
```

```
Based on a casual disassembly of MSCDEX (which I did on the urging of
Brian Livingston, who correctly suggested there was something fishy
going on here), it is now clear to me that DRVSET doesn't truly work
for unloading MSCDEX, because MSCDEX "patches" DOS. A true NOMSCDEX
would back out these patches too. Or do the patches occur
only when running MSCDEX /S on a server, so that using DRVSET for
NOMSCDEX "would" work except when running MSCDEX /S? Anyway,
programs like MSCDEX do enough gross low-level things to DOS that
trying to back out the changes and unload them is a risky
proposition at best.
```

```

*,
#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <ctype.h>
#include "currdir.h"

void fail(char *s) { puts(s); exit(1); }

main(int argc, char *argv[])
{
    char *cmd;
    CDS far *drv;
    int drive, drvoff, i;

    /* to just turn off drives, rename program DRVOFF */
    drvoff = (strstr(argv[0], "DRVOFF") != 0);

    /* what drive do they want (accepts letters and numbers) */
    if (argc < 2) fail(drvoff?
        usage drvoff [drive]:
        usage drvset [drive] NET PHYS REDIR SUBST JOIN OFF COPY [drv2] "
        "BACKOFF [ofs]");

    drive = toupper(argv[1][0]) - 'A';
    if (! (drv = currdir(drive)))
        fail("can't get current directory structure");
    if (drvoff) { drv->flags = 0; return 0; } /* just turn off drive */
    for (i=2; i<argc; i++)
    {
        cmd = strdup(argv[i]);
        if (strstr(cmd, "COPY")) /* COPY one CDS entry to another */
        {
            CDS far *drv2 = currdir(toupper(argv[i+1][0]) - 'A');
            if (! drv2) fail("can't copy from invalid drive ");
            if (memcmp(drv, drv2, (sizeof *drv) * 0x5B - 0x51),
                0)
                continue;
        }
        else if (strstr(cmd, "BACKOFF")) /* set backweek offset */
        {
            drv->backslash_offset = atoi(argv[i+1]); i++;
            continue;
        }
        /* COPY and BACKOFF let you simulate SUBST (Big deal!) */
        /* change drive attributes */
        if (strstr(cmd, "OFF")) drv->flags = 0;
        if (strstr(cmd, "CDROM")) drv->flags |= REDIR_NOT_NET;
        if (strstr(cmd, "REDIR")) drv->flags |= REDIR_NOT_NET;
        if (strstr(cmd, "NET")) drv->flags |= NETWORK;
        if (strstr(cmd, "SUBST")) drv->flags |= SUBST;
        if (strstr(cmd, "JOIN")) drv->flags |= JOIN;
        if (strstr(cmd, "PHYS")) drv->flags |= PHYSICAL;
    }

    /* print current drive state */
    if (! drv->flags) fputs("INVALID ", stdout);
    if (drv->flags & REDIR_NOT_NET) fputs("REDIR ", stdout); // CDROM
    if (drv->flags & NETWORK) fputs("NET ", stdout);
    if (drv->flags & SUBST) fputs("SUBST ", stdout);
    if (drv->flags & JOIN) fputs("JOIN ", stdout);
    if (drv->flags & PHYSICAL) fputs("PHYSICAL ", stdout);
    putchar('\n');
    return 0;
}

```

DRUSEL also has an option to copy from one CDS entry to another and an option to set the CDS backslash. It's not an unnecessary and somewhat bogus but perhaps intriguing way this simulates the effect of the SUBST command.

```
D:\UNDOC2>dir e
Invalid drive specification
D:\UNDOC2>drvsset e, copy d subst backoff 9
SUBST PHYSICAL
D:\UNDOC2>dir e

Volume in drive E has no label
Volume Serial Number is 1B51-11E7
Directory of E:\

.                <DIR>          02-16-93  11 52a
..               <DIR>          02-16-93  11 52a
NAMCLUST.EXE    9812 07-14-93  7 31p
CLUSTNAM.EXE   8102 07-10-93 10 42p
CHAPB.TXT      321850 07-20-93  8 49a
? ... etc .
```

This is close enough to the actual effect of SUBST that an examination of the code in DRUSEL.C should give you some idea of how the real SUBST command must work.

```
CDS far *dest = currdir(*E' - 'A'),
CDS far *src = currdir(*B' - 'A'),
_fmemcpy(dest, src, _osmajor >= 4) ? 0x58 : 0x51),
dest > flags SUBST,
dest->backslash_offset += strlen(src->currdir),
```

As another example of modifying the CDS, see the MASTDRV program in Chapter 2, which changes the location and size of the CDS—that is, the FANSDRIVE value—of the FV.

For more on file redirection, about drives and directories, to get yourself another cup of coffee and another slice of pizza, and proceed to the next installment in our saga of the DOS file system. It's now time to look at the data structures that DOS uses when you open a file.

### System File Tables (SFTs) and Job File Table (JFT)

The System File Tables are the backbone of the DOS file system and have been present in DOS since version 2.0 when Microsoft introduced a revised file operations. Before that, DOS used File Control Blocks (FCBs), which are discussed briefly later in this chapter. An SFT contains the state of an open file. It includes associated information with a directory entry and with a cluster on disk, keeping track of the current position within the file, the position determining current file size, and managing file reads and data stream operations as they are needed. An information contained in each directory entry for the file is then taken from the SFT when the file is closed; it is brought back into the SFT when the file is opened.

When a program opens a file using DOS function 31h (Open File) or 6Ch (extended Open Create), it gets a file handle that are related to the file when reading with function 3Fh (writing with function 40h) on the file. The file handle is what Microsoft calls a "magic cookie," meaning that you can use this handle to access the object—in this case, a file—without assigning any particular meaning to the handle. For example,

```
#include <io.h>
// ...
char buf[12];
int f = open("foo.bar", O_RDONLY);
if (f == -1) fail("couldn't open file");
read(f, buf, 12);
close(f);
```

After checking that  $1 = 1$  (that is, that the file was successfully opened), most applications treat file handles as "use or throw." An action such as `open()` returns the handle to you, and you faithfully pass copies of file handles such as `read()`, `write()`, and `close()`. With the exception of the predefined handles `_stdout`, `_stdin`, `_stdev`, and `_timer` (see 3.1), and some (see 3.1, `stdarg` and 4.1), standard applications generally attach to file handles through the common use of a file handle. DOS, however, certainly does attach to file handles through the variable, and it's somewhat unusual for an application to do so too.

A DOS file handle starts with a file system (FAT) 21h function, 33h and 64h, a simple an array of 256 characters, the File File Table (FIF) is an array of BFFs, indexed by file handles, and handling SEI indices. In other words,

```
int f = open(...);
sft_index = jft(f);
sft_entry = sft(sft_index);

...

sft_entry = sft(jft(f));
```

A file pointer is a pointer to a Non-1 NetWare versions 2 and 3 which, as discussed in Chapter 4, assign `root` to the root of the local or network server. SEI uses this as a reverse index into its own file system, and the file system is possible to SEI.

Address accesses don't make a file SEI. A far pointer to the last SEI in the chain can be found at `file` in `SEI`. Each DOS process has its own SEI. A far pointer to a process's SEI can be found at other 34h in its PSP (see Chapter 7).

## FILE\* vs. File Handles

Before we go any further, we need to distinguish DOS file handles from the file identifiers used in compiler run-time libraries. Many DOS applications use a compiler run-time library rather than calling directly down to DOS. In some cases, like the C `open()` function declared in `IO.H`, the DOS version of the C run-time library usually does happen to return a DOS file handle. However, in other cases, such as the C `fopen()` function declared in `STDIO.H`, something completely different, `vertained fopen()` returns a pointer to a FILE structure (FILE\*).

FILE is a structure declared in `STDIO.H`. It is an indication of the success of the C run-time library that most programmers think purely of FILE\* without having to give any consideration, and perhaps without ever having looked at FILE itself. The C run-time library declares a function `fun` which, given a FILE\*, can return the corresponding file handle that the programmer would have had it called `open()` rather than `fopen()`. In most C compilers for the PC, the FILE structure contains a field for the DOS file handle, and `fileno()` simply returns the value of this field.

```
#include <stdio.h>
//
FILE *f = fopen(...);
sft_entry = sft(jft(fileno(f)));
```

The superficial similarity between FILE\* and file handles is a surprisingly large source of confusion. The technical support department for one vendor of programmer's products reports the confusion between FILE\* and a DOS file handle as one of the most frequently occurring problems among its (supposedly quite sophisticated) clientele. This confusion is particularly rampant among those trying to increase the number of available file handles. They call INT 21h AH=67h to get more DOS file handles and then wonder why `fopen()` still fails after twenty `open()`s. We'll get to this topic later in this horrendously long chapter.

All DOS systems have at least one SEI entry; many have 20 or more. Just as the LASTDRIVE value in CONFIG.SYS sets the size of the DOS FILES\_ establishes the number of SEI entries. (Incidentally, this is a well-known answer to an end user's question: "What exactly does the FILES directive in CONFIG.SYS do?" see Jeff Proise, *What THE 86 Does: PC Maintenance*, November 2, 1991. FILES\_ is a table of pointers to such commands present in CONFIG.SYS. Every file handle that a program opens flows from DOS1\_acts\_ through the FILE\_ to one of the SEIs.

When you ask DOS to open a file by calling functions 3Bh or 6C\_ you are calling a higher-level function like 1\_ to open a file. It internally calls one of the DOS functions for you; the tool wrapper takes place.

First, DOS uses the current PSP (see Chapter 6) to locate your FILE\_ actually, the FILE\_ of whatever PSP happens to be current. But we assume for this discussion that the current PSP belongs to your program. (I ought not see Chapter 9 on IRRs.) DOS searches through the current PSP's FILE\_ to find a slot that is not currently in use, and remembers the index into the table for the first such free slot that it finds. This index into the FILE\_ eventually becomes the handle associated with the open file, assuming of course that DOS searches for a free FILE\_ slot first because a free FILE\_ slot if DOS can't open a file and doesn't need to do anything more. DOS would return to you with an error code 4 (too many open files). A simple heuristic to increase the size of its FILE\_ is to set on 67h. See Maximum Handle Count, discussed later in this chapter.

In its next attempt to find a free FILE\_ entry, DOS will search the chain of SEIs, looking for the first SEI entry that is available for use. If no suitable entry is found, again DOS fails the request, and returns error code 4. Factors SEI\_ and situation\_ which is normally set to string \$drive\_ from FILE\_ call\_ calling to set on 67h to increase the FILE\_ size does not help; the best bet is to raise FILES\_ in CONFIG.SYS and reboot, though you will see later that there is another way to get more SEI entries.

If both a free handle\_ FILE\_ entry\_ and a free SEI\_ entry\_ exist, DOS calls the SHARE\_ file\_ open hook function; see the discussion of SHARE\_ later in this chapter. Assuming SHARE\_ allows the file open to succeed, DOS determines the drive on the filename you asked to open and uses this to index into the CDS\_ to locate the proper CDS\_ entry for the drive where the requested file resides. If you passed a relative path name such as "C:\OOBAR", DOS uses the current directory in the CDS\_ if installed; the couple switches ASSIGN\_ APPEND\_ and LASTOPEN\_ do bear some.

From the CDS\_ DOS now determines if the file you want to open is on a network drive. If so, DOS issues the appropriate IN\_ 2Bh AH\_ heads to have any installed network redirectors open the file. As noted at the end of the section on the network redirector interface, when DOS calls one of the Open or Create redirector functions (IN\_ 2Bh AH\_ 116h\_ AV\_ 117h\_ or AV\_ 112h\_), it passes the address of a free slot in the SEI\_ entry\_ which it expects the redirector to fill.

If you are calling with a path name drive\_ DOS always has CDS\_ to extract a pointer to be File's DPR. DOS uses the DPR to locate the entry's root device; it expects then the necessary information to convert a user's filename to LFNs and to construct the block device driver that handles the driver. From the device driver's header, DOS locates function pointers to the driver's Interrupt and Strategy routines, which are used to communicate to and from the drive.

Armed with this information, using the DPR, the DOS kernel calls the device driver to read the driver's root directory into one of the DOS buffers, unless of course the root directory is already low, because of the buffers recalled from the BUFFER program. Listing 8.8.1 at the end of the book often include directory and FAT sectors.

If the supplied path contains any subdirectories, the DOS kernel searches the root directory of the drive, trying to locate the first component of the supplied path (e.g. the "VENDOR" in "C:\VENDOR\2\FILES\OOBAR"). If it isn't found, the function fails with "Path not found" (error code 3).

If DOS adds this top-level directory, however, it then converts the starting cluster value in its directory entry to an LFN. It then passes the LFN to the device driver, which reads that subdirectory into an other buffer or uses it to locate the subdirectory if it's already in the buffers. This process continues if DOS has traversed all directories in the supplied path string, and only the file name remains to be read.

As this is quite similar to the code shown earlier in NAMELUST.C (Listing 8.5), which turns a file or directory name into a starting cluster number. For example:

```
C:\>DIRDOC2\CHAP8>namclust foo.bar
```

NAMELUST first calls filename() to convert "foo.bar" into "C:\DIRDOC2\CHAP8\FOO.BAR", and then calls get\_drive() to give the information necessary to read in its root directory. After reading in the root directory, it searches for "DIRDOC2", just as MS-DOS itself would if we issued an INT 21h AH=3Dh to open that file.

```
C:\>namclust undoc2
```

```
C:\>DIRDOC2 ==> 48 (sect 973, 0x00003ed)
```

After seeing DIRDOC2's starting cluster number, as found in its entry in the root directory, DOS won't read in this sub-directory list, of course. DOS first checks the buffers; DOS would next search for "DIRDOC2" subdirectory for "CHAP8". If CHAP8 is not found in the first cluster of this sub-directory, DOS looks for the possible next cluster for the subdirectory and searches there too. Again, this is similar to the operation of NAMELUST.

```
C:\>DIRDOC2>namclust chap8
```

```
C:\>DIRDOC2\CHAP8 ==> 51 (sect 1021, 0x00003fd)
```

Assuming CHAP8 is found, DOS gets its starting cluster number and reads in that cluster. It is now at the final level in the file name: "FOO.BAR".

At this point, and no earlier, DOS determines whether you are dealing with a real file or with a named character device, such as CON or LPT1. It does so by searching the list of installed device drivers (see the device driver chapter's "kernel program" chapter). If DOS finds an exact match for the file's extension of the pathspec you gave it, it ignores any extension in this test; it opens the device rather than a file. This means that all the named devices seem to exist in all directories of the file system. They also exist in a searching sub-directory named DIR, even though DIR.DIV fails if you attempt to open a non-openable regular disk of extension with the same name as one of the installed devices. This is why, for example, READ would be a very bad name for a device installing this device would prevent you from accessing a README file. This is one reason device names select common words tend to exclude dollar signs, for example "CLOCK\$".

If an installed device name is found, DOS searches the last directory for the filename and the extension. Again, this can be simulated with NAMELUST:

```
C:\>DIRDOC2\CHAP8>namclust foo.bar
```

```
C:\>DIRDOC2\CHAP8\FOO.BAR ==> 5576 (sect 89421, 0x00015d4d)
```

Assuming the file or device exists, DOS sets the pointer to the first free SEF-located character byte and the index of this SEF entry system somewhere in the program's HPI. DOS will later return the "somewhere" part of the call to program as the file handle. If DOS can't match the name, it instead returns error code 1 ("file not found").

What happens next depends on whether you're opening a file or a device. For a file, DOS opens a pointer to the root of the file's directory entry into the corresponding fields of the SEF entry. The root location includes the starting cluster for the file. DOS also copies in a pointer to the DPH for the drive on which the file resides; the DPH in turn points to the device driver header for this drive. DOS also sets the file pointer field of the SEF to zero, indicating the beginning of the file. For a device, the SEF entry includes a pointer to the device driver header.



For both files and devices, DOS then returns to you with the AX register set to the handle value that it reserves at the start of the process. You now refer to the file using that handle, the major `create` and `write` system calls. This file handle is merely an index into the current PSP's HLT, and the byte at that index, HLT[handle], is itself an index into the SEI.

If you're creating a new file rather than opening an existing one, this same basic sequence of events is followed. The only significant difference occurs when the SEI entry has been filled out before the DOS file creation comes back to you with a handle. The new directory entry for the just-created file is put back into the DOS buffer for that PDS, and the directory for that buffer, via the HLT[ERFC] entry, is set to zero. This is the DOS's way of writing the buffer out to disk as soon as possible, certainly before reassigning it. A directory entry is created immediately for your new file, though with a length of zero.

Each time you read from or write to the file, referencing it by means of the handle, DOS uses the supplied handle to index into the HLT associated with the current PSP, and it uses the value it finds there to index into the SEI. The handle is then used to perform the requested operation on the file or device, as indexed by its corresponding SEI entry. The file pointer and date timestamps in the SEI are updated accordingly. Data transfers also normally involve the DOS buffer chain. To summarize:

current PSP → HLT → SEI → buffers, DPB, device driver

When you close a file, DOS accesses its SEI just as for reading or writing. If the file has been written to, as indicated by a status bit in the SEI attribute word, DOS updates its directory entry with information from the SEI, which reflects the latest size, time, and date. Also, any dirty buffer is flushed to disk. If the file has not been written to, DOS skips these steps.

DOS's directory entry handle count table of the SEI entry, for quickest access, reflects the fact that any handle is being directed from the SEI. The SEI index in the HLT handle table is replaced by the value HLT[handle] which indicates an indirect, variable slot. Because HLT[255] indicates any HLT entry can contain a valid SEI index, HLT[254] therefore, by maximum possible value, you can set. However, remember that Novell NetWare uses negative file handles, as discussed in Chapter 4, setting a very large HLT[255] value can reduce the number of files you can open on NetWare servers.

If the newly free HLT entry is the only one using the SEI entry, decrementing the handle count brings it back to 0 and makes the entry available for reuse the next time someone calls DOS Open or Create. Some programs keep multiple files open with handles, but despite many programs' insistence on having 30 to 50 files available, it is rare for the number of SEI's in use to grow much larger than 25 or 20. As an illustration with the HLT program a little later on, when you look at the SEI there usually isn't much to see.

## How Many FILES?

Earlier, we walked through the DOS disk buffers to determine the value of BUFILES—we can likewise walk through the SEI's to determine the value of FILES. Just as with BUFILES, this value is normally set to CONFIG.SYS, although you can alter it on the fly with a utility like Quarterdeck's FILES.COM.

SEI.WALK.COM, Listing 8-17, determines the value of FILES, and prints out the address and size of each SEI. We'll reach size and code in Chapter 9, see Listing 3-8, which showed how to port this read-only code to protected mode Windows. The result was a second version of SEI.WALK.COM, Listing 8-18, that used DOS and/or Windows API calls. Here, we focus on the SEI entries.

SEI.WALK determines the FILES value by threading through the SEI entries and keeping count of the entries in each table. The first SEI always holds the possible open entries if its assembly is into MSDOS.SYS. If FILES=40 appears in CONFIG.SYS, for example, then DOS allocates a second SEI, etc., enough for 35 more files, and chains it to the first SEI. Since each

header consists of a count of the number in its associated table, together with a pointer to the next header, it's easy to count the possible files.

If DO\_FCBS is enabled, SFTWALK can also determine the value of FCBS\* by walking the chain of System FCBS. These have the same format as the SEFs. We briefly discuss System FCBS later in this chapter.

### Listing 8-17: SFTWALK.C

```

/*
SFTWALK.C -- Count FILES* by Walking SFTs
for protected mode Windows version, see \UNDOC2\CHAP3\SFTWALK.C
Andrew Schulman, March 1993
from Undocumented DOS, 2nd edition (Addison-Wesley, 1993)
*/

#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <dos.h>

typedef unsigned char BYTE;
typedef unsigned short WORD;

typedef struct _sft {
    struct _sft far *next;
    WORD num;
    // the actual SFT entries start here
} SFT;

int sftwalk(int offset, char *name)
{
    static BYTE far *sysvars = (BYTE far *) 0,
    SFT far *sft;
    int files = 0;
    if (! sysvars)
    {
        _asm mov eh, 52h
        _asm int 21h
        _asm mov word ptr sysvars+2, es
        _asm mov word ptr sysvars, bx
        if (! sysvars) return 0;
    }
    sft = *((SFT far * far *) &sysvars[offset]);
    while (FP_OFF(sft) != 0xFFFF)
    {
        files += sft->num;
        printf("%s at %p    %u files\n", name, sft, sft->num);
        sft = sft->next;
    }
    return files;
}

main()
{
    int files, fcbs;
    files = sftwalk(4, "SFT");
    printf("FILES=%d\n\n", files);
#ifdef DO_FCBS
    fcbs = sftwalk(0x1A, "System FCB table");
    printf("FCBS=%d\n", fcbs);
#endif
    return files;
}

```

Even if Windows or a program such as QEMM's FILES.COM has added an SFT (and possibly loaded it into upper memory), SFTWALK finds it. For example,

```

C:\WINDOWS\CHAP8>sfwalk
SFT @ 0116:0000 -- 5 files
SFT @ 1014:0000 -- 35 files
FILES=40

C:\WINDOWS\CHAP8>\qemm\files +10
FILES=40 before
10 files added
FILES=50 now

C:\WINDOWS\CHAP8>sfwalk
SFT @ 0116:0000 -- 5 files
SFT @ 1014:0000 -- 35 files
SFT @ 1838:0000 -- 10 files
FILES=50

```

From the SFWALK output, you can see that the earlier mention of sfwalk was an oversimplification. Because there is a chain of several SFIs rather than a single array, indexing into the SFI requires walking the linked list to first find the correct table. For example, sfwalk 1014:0000 in the above configuration sfwalk would be somewhere in the second SFI table that starts at 1014:0000.

But what if you display the size of each SFI entry. We've talked a bit about how DOS uses SFIs and about how to walk the SFI list, but we haven't yet seen what an SFI entry looks like. Luckily, the SFI is mostly opaque, so there's not much to look at, but we can create a fake SFI by running Windows, which keeps many files open. After running SFWALK within a Windows DOS box to get the address of each SFI in the SFI chain, we can then use DIBUG to look at one of the SFIs.

```

C:\WINDOWS\CHAP8>debug sfwalk.exe
-g
SFI @ 0116:0000 -- 5 files
SFT @ 1014:0000 -- 35 files
SFT @ 1838:0000 -- 10 files
FILES=50

Program terminated normally
-d 1014:0000
1014 0000 00 00 19 18 23 00 01 00 00 00 20 42 10 AA 13 16          # B
1014 0010 01 F9 08 40 19 6A 18 C0 14 00 00 00 02 00 00 00    # 2
1014 0020 00 F1 01 00 00 05 45 47 41 38 30 57 4F 41 46 4F  .....EGAR0W0AFO
1014 0030 4E 00 00 00 00 01 00 09 1F 00 00 F9 08 00 00 00  N.....
1014 0040 00 01 00 00 00 20 42 10 AA 13 16 01 F7 08 40 19  .. B .....
1014 0050 6A 18 80 20 00 00 A0 1E 00 00 00 00 F1 01 00 00  j.....
1014 0060 02 45 47 41 3A 30 57 4F 41 46 4F 4E 00 00 00 00  .EGAR0W0AFOF.....
1014 0070 01 00 09 1F 00 00 F7 08 00 00 00 00 01 00 00 00  ..

```

What else we get here. Each SFI starts off with a header shown in SFWALK as DW0020 pointing to the next SFI. The offset 0FFFF indicates the end of the chain, followed by WORD with the number of files in this SFI, followed by SFI entries themselves. In the SFI at 1014:0000 the first four bytes 00000019 indicate that the next SFI is at 1838:0000, the next two bytes 2300 indicate that there are 23 files, just as SFWALK claimed.

SFWALK does a good job of looking at the SFI entries themselves, but from the DIBUG core dump you can plainly see the file names EGAR0W0AFO and EGAR0W0AFOF. Notice that the SFI includes only the file name and system name, not the full path name. However, the SFI also includes the starting cluster for each file, so we can know from the CLUSTNAME program, if necessary, exactly in this starting cluster into the full path name.

One could determine the size of a SFI entry simply by subtracting the start of one file name (01F7 hex here) from the start of the next, for example offset EGAR0W0AFOF - offset EGAR0W0AFO = 0x000026h = 38h bytes. As crazy as it sounds, there is code in Windows

from here, just this: see "CON CON CON CON CON" in Chapter 1. Programs that want to deal with SFTs in the widest range of DOS versions have to somehow deal with the fact that the SFT size and offset differs in DOS 3.0, DOS 3.0, and DOS 3.1, and of course may change again in future DOS versions, though so far it has remained stable from DOS 3.1 through DOS 6.0.

In DOS 3.x and higher, each SFT entry is 300h bytes, and the filename is located at offset 200h in each entry. We use this information in our next program, which also takes us back to look more closely at the JFT and its relation to the PSP.

### Filename From Handle

Sometimes you need to know the name of a file and have only its handle available. One important example of this is when you use the DOS redirection facility to redirect stdout to a file rather than to CON, its normal destination. While the stdout handle is always 1, there's no documented way of telling from a shell or running program the *name* of the file to which this (or any other) handle corresponds.

This next program, H2NAME.C (Listing 8-18), converts file handles to names by combining several macros created DOS for us. It consists primarily of the function `h2name`, which when passed a PSP and a handle, returns the filename to which that PSP's handle corresponds on current system.

We'll use `h2name` just like the function `getname` and as in other programs, H2NAME.C also includes a standard DOS-style H2NAME with a PSP number on the command line; you can get the PSP number via redirection to process 0 or the COMMAND program in C:\\_cmd\_. H2NAME.C increments a special register to that process. Otherwise, it concatenates all open files belonging to H2NAME's process, the `_psp_global` found in most C compilers for the PC. To see anything interesting, it is, as you should expect, H2NAME's output to a file so that this redirected output file shows just the open file information. H2NAME sends its output to stderr, so the output still stays by you, at least. Here, the program's unredir'd output:

```
C:\VHD0C2\CHAPP8>foo bar < h2name c
Files for 7976
0 -> 5 -> H2NAME.C
1 -> 4 -> FOO BAR
2 -> 3 -> CON
3 -> 0 -> AUX
4 -> 7 -> PRN
*
```

Standard output still is always file handle 0. Here, we can see that `stdin` has been redirected from H2NAME.C's window to SFT entry 3. Standard output, `stdout`, is always file handle 1, and again you can see that H2NAME knows its standard output has been redirected to the file FOO BAR, with SFT entry 4. We took a different redirect `stderr`, so this file handle 2 shows up in the H2NAME output as CON.

Only the name and extension are reported. Getting a complete pathname would require using the CLUSTNAM code (Listing 8-6) to process the starting cluster number stored in the SFT.

### Listing 8-18. H2NAME.C

```
/*
H2NAME.C -- Convert file handle into file name.
Jim Kyte, 1991, Revised Andrew Schulman, July 1993
The file handle is a JFT index. JFT[handle] is an SFT index. The
SFT contains information on the file, including the name. However,
H2NAME only gets the name and extension. To get the full pathname,
convert cluster number in SFT entry to a pathname (see code in
CLUSTNAM.C to see how to turn clusters numbers into full pathnames).
*/
#include <stdio.h>
```

```

#include <stdio.h>
#include <string.h>
#include <dos.h>
#ifdef __JRB00__
#include <mem.h>
#else
#include <memory.h>
#endif

#ifdef MK_FP
#define MK_FP(s,o) ((void far *)
    (((unsigned long)(s) << 16) | (unsigned)(o)))
#endif

char * h2name( unsigned psp, int h, int *sft_handle )
{
    static char name[15], /* will hold file's name */
    static unsigned far *sft_ptr = (unsigned far *) 0;
    static unsigned nrofs,
    static int sfts ze,
    unsigned far *ptr,
    signed char far *sptr, far *htbl;
    int sftn,
    memset( name, 0, 15 ), /* blank out the static name */

    /* create pointer to handle table (JFT) */
    htbl = *((char far * far *) MK_FP( psp, 0x34 ));

    if ( ! sft_ptr ) /* one-time initialization to get SFT info */
    {
        __asm mov ah, 52h
        __asm int 21h
        __asm les bx, dword ptr es:[bx+4] /* SFT chain = SysVars[4] */
        __asm mov word ptr sft_ptr+2, es
        __asm mov word ptr sft_ptr, bx
        switch( __osmajor )
        {
            case 2: sftsize = 0x28, nrofs = 4; break;
            case 3: sftsize = 0x35, nrofs ( __osminor == 0 ) ? 0x21 : 0x20, break;
            default: sftsize = 0x3B, nrofs = 0x20; break; /* DOS 4.0, 5.0, 6.0 */
        }
    }

    ptr = sft_ptr;
    *sft_handle = htbl[h],
    if ( htbl[h] >= 0 ) /* now if handle is valid */
    {
        sftn = htbl[h]; /* get index into SFT list */
        while ( FP_OFF(ptr) != 0xFFFF )
        {
            if ( ptr[Z] > sftn ) /* then target is here */
            {
                sptr = (unsigned char far *) &ptr[3];
                while ( sftn-- ) /* so skip down to it */
                    sptr += sftsize;
                memcpy( name, &sptr[nrofs], 11 );
                return name; /* found and copied; done */
            }
            sftn -= ptr[Z]; /* not here; reduce index */
            ptr = (unsigned int far *) MK_FP( ptr[E], ptr[Q] ),
        }
    }
    strcpy( name, "UNKNOWN" ),
    return name; /* reached only by error */
}

void fail(const char *s) { puts(s); exit(1); }

main( int argc, char *argv[] )
{
    char *s,

```

```

unsigned psp;
signed int max_files, sft_handle, jft_handle;

if (argc < 2)
    psp = _psp, /* display files for this program */
else
    sscanf(argv[1], "%i", &psp), /* take PSP from command line */

max_files = (osmajor > 3) ? ((int) far *) MK_FP(psp, 0x32) : 20;

if (*(unsigned far *) MK_FP(psp, 0) != 0x20C0) /* check for INT 20h */
    fail("that's not a PSP!");

fprintf(stderr, "Files for %04X\n", psp);

for (jft_handle = 0, jft_handle = max_files, jft_handle++)
{
    s = h2name(psp, jft_handle, &sft_handle);
    if (sft_handle < 0)
        fprintf(stderr, "%2d ==> %2d\n",
                jft_handle, sft_handle), // unused, or Novell
    else
        fprintf(stderr, "%2d ==> %2d ==> %s\n",
                jft_handle, sft_handle, s);
}

return 0;
}

```

H2NAME's h2name works. The PSP contains a pointer to the HFT—and usually the HFT itself. The h2name function first creates a ptr pointer to the handle table. H2A for the specified PSP. It then checks the SEI size and SEI offset to set up the SEI pointer and two variables: stride and offset. Both SEI size and the offset within the SEI of the file or device name, based on the DOS version in use.

With these preparations out of the way, h2name uses the supplied handle value as an index into the HFT. If the value found there is non-negative, indicating a valid handle, it is an SEI index. Recall that negative file handles have special meaning in Novell NetWare.

It then goes on to work through the linked list of SEIs until it finds the SEI containing the desired file or device name. sft = ptr + sft \* stride. Next, this program's admittedly odd way of saving sft > max\_files. Each time a SEI is skipped, sft = SEI - the number of entries in the skipped block is subtracted from the desired index. ptr = ptr + 2 \* sft. The index is always relative to the current block rather than to the absolute beginning of the SEI linked list.

When the correct block's offset pointer is set to the first byte of its first SEI entry (sptr = &ptr[sft \* offset + SEI\_offset] + a skipped sft \* stride), determining the index each time, until we hit bytes hex zero. When this happens, the SEI entry under the pointer is the one you're looking for. The index field offset value is then added to sptr, the chosen bytes at the resulting location are copied into the static buffer, and the program returns a pointer to the first byte of the buffer.

Clearly, we could apply the same technique to other attributes found in the SEI, h2attr(), for example, would return the file attributes rather than the filename, and h2cluster() will give you the disk.

H2NAME's h2name() for each HFT entry. H2NAME's ability to look at the HFT for any PSP specified on its command line will come in handy later when we need to test the HANDLE and FILES programs. If H2NAME were built as a Windows program, using the techniques shown in FILES (Listing 8.19) and explained further in Chapter 5, it could reveal the HFT for KRNL386 or any other program running in the System VME. As Chapter 5 noted, Windows programs have PSPs, just as DOS programs do.

### What Files Are Now Open?

Our next utility, `FILES`, displays information about all open files and devices on your system. Normally, there aren't many open files in the `FILES` to see anything other than `ALX CON PRN`, an `IPRN` redirect, `FILES` output to a file, for example, `files > files.log`, or an unnecessarily redirected input from a file: `files < foo.bar`.

```
D:\MUNDCC2>files > files.log < foo.bar
```

```
D:\MUNDCC2>type files.log
```

#	Filename	Size	Attr	cRef	Owner	Cluster, DB, BPB
[SFT @ 0116:00CC -- 5 files]						
0	AUX	0	0000	7	9077	DEV 0070:0035 0
1	CON	0	0000	19	9077	DEV 0070:0023 0
2	PRN	0	0000	7	9077	DEV 0070:0047 0
3	FOO.BAR	7391	0020	2	6AC2	D 29444 0PB 0FF1:0000 0
4	FILES.LOG	0	0020	2	6AC7	D 0 0PB 0FF1:0000 0
[SFT @ 1014:0000 -- 35 files]						
[SFT @ 1858:0000 -- 10 files]						

When any input is captured redirected, `FILES` inherits an open file from `COMMAND.COM`. Note how `FOO.BAR` and `FILES.LOG` each have a reference count of 2.

You can also use `CLUSTNAM` to confirm the cluster numbers `FILES` displays. Only one file in the `FILES` `FOO.BAR` has a cluster number `29444`, to which output from the `FILES` program was redirected, as in the file of course just a directory entry with no associated data clusters.

```
D:\MUNDCC2>clustnam d: 29444
```

```
D:\MUNDCC2>FOO.BAR
```

Besides a file check, to be sure that the output from `FILES` is accurate, this also shows that the `CLUSTNAM` code (Listing 8.5) can be showed into `FILES` to output the full pathname for each file entry with an associated data cluster. The same point was made earlier regarding the `FILENAM` program, a shortcoming of the `CLUSTNAM` code that may have implications for you.

As you'll see in `FILES` (Listing 8.19), the owner `FILES` displays is a pointer to the `FILE` object, via a `PSN` and, correspondingly, `FILES` will use `UDMM` code (see Chapter 7) to turn the owner's `PSN` address into a readable name.

However, these values are not always legitimate `PSN` addresses. The `ALX`, `CON`, and `PRN` entries show a cluster of `0077`. Running `UDMM` indicates that this is not a valid `PSN`. Instead, the value is apparently the cluster's `PSN` at the time that the `SYSTEM` initialized on `IO.SYS` or `IBMBIO.COM` loaded them. `SYSTEM` relocates itself to the top of memory, accounting for the high address.

As a more interesting example, consider the output from `FILES` when running in a DOS box under Windows 3.1 Enhanced mode:

```
D:\MUNDCC2>files > ps
```

#	Filename	Size	Attr	cRef	Owner	Cluster, DB, BPB
[SFT @ 0116:00CC -- 5 files]						
0	AUX	0	0000	17	9077	DEV 0070:0035 0
1	CON	0	0000	51	9077	DEV 0070:0023 0
2	PRN	0	0000	17	9077	DEV 0070:0047 0
3	WIN386.SWP	1769472	0020	1	1AAE	C 6506 0PB 0116:13A4 0
4	CONRF.FOH	25408	0020	1	1F00	C 2513 0PB 0116:13A4 1
[SFT @ 1014:0000 -- 35 files]						
6	SH.EXE	53792	0020	1	1F00	C 1218 0PB 0116:13A4 1
7	VGAFIX.FOH	5360	0020	1	1F00	C 2510 0PB 0116:13A4 1
8	GDI.EXE	275264	0020	1	1F00	C 3057 0PB 0116:13A4 1

```

0  SSERIFE  F0W      64544  0020    1  1FD0    C  2547  DPB  0116:13A4  1
10 USER     EXE      538406  0020    1  1FD0    C  3072  DPB  0116:13A4  1
11 WINFILE  .EXE     146864  0020    3  1FD0    C  1794  DPB  0116:13A4  1
12 KEYBOARD DRV     7568  0020    1  1FD0    C  2301  DPB  0116:13A4  1
13 VGA      .DRV     73200  0020    3  1FD0    C  2204  DPB  0116:13A4  1
14 COMPT    DRV     9280  0020    1  1FD0    C  2253  DPB  0116:13A4  1
15 FILES    EXE     14698  0020    1  1FD0    C  12643 DPB  0116:13A4  2
16 ANTQUA   TTX     59776  0020    1  1FD0    C  2947  DPB  0116:13A4  1

```

[SFT @ 1B19 0000 -- 30 files]

Note WIN386.SWP, the SFT entry #3 in the output above. As its name suggests, this is the Windows 386 Swap File, a file swap file which Windows uses for draggaging when you don't have a permanent swap file. The SFT entry #115 displayed indicates WIN386.SWP's owner is 1AA2h. We can run `FILENAME -> 1AA2` to make sure this all makes sense.

```

D:\>UNDOC2>h2name 1aa2
Files for 1AA2
0  > 1  > CON
1  > 1  > CON
2  > 1  > CON
3  > 0  > ALX
4  > 2  >> PRN
5  > 3  >> WIN386 SWP
6  > 1

```

Note the path `FILE\FILE\PSN\1AA2h` has a file handle that points back to this SFT entry. Running `FILENAME -> 1AA2` or `MEM -> P` shows that PSN 1AA2h corresponds to WIN386.EXE.

```

D:\>UNDOC2>udname
1AA2 1AAE 0452 ( 17696) Env at 1AA2 C:\WIN31\sys\win386.exe

```

We note the owner ID 0B2h of SFT entries 4-16 in the FILE output. Specifying the number 1101h or 112NAME shows that the process processes the message "that's not a PSN". 112NAME merely checks the first two bytes of the specific segment to see they are the 0x201D signature which are the opcode bytes for the `INSTR 201` instruction that begins every PSN.

The problem is that when running 112NAME from within a DOS box, the last column in the FILE output is the virtual machine VM ID of the file system. 0 indicates files or devices opened by the Windows static file driver, the System VM, where Windows applications run. 2 indicates the first DOS box in the first FILE system, and in a first DOS box VM 2. As indicated by the last column of FILE output, because 1101h is only a valid PSN within another virtual machine, VM 1. In fact, 0 is the virtual machine is the System VM in your Windows applications are run and this PSN's address space is that of the System VM's address space. Even if the DOS box happened to have a PSN 1101h, it is still the only legitimate seek in Chapter 6. As discussed in Chapter 3, view `FILE` a PSN of the System VM carries a DOS program run from within WINSTART.BAT, 2 a protected-mode DOS program which gets its memory from other VMs or simply 3 a Windows application. When viewing the `WINPSN` program from Chapter 3, PSN 1101h or the equivalent on your machine shows up plain as day.

```
1F00 7FF5 KRNL386 5196 0000 (52 files, 18 open)
```

But if the SFT entries for `COLORATION`, `SHREVE`, `AGSMITH`, and so on all belong to a PSN that is only visible within the System VM, why then are these open files located in the global SFT variable to all processes in all VMs. In other words, why doesn't Windows declare the SFT as instance data because while Windows instances other DOS data structures such as the CDS on a per-VM basis,



Windows can't instance the SFT. How else would SHARE work? See the SHARE discussion later in this chapter.

On the other hand, parts of SFTs can be allocated on a per-VM basis. The Windows SYSTEMINFO has a PerVMFiles setting that controls how many SFT entries each VM can add to its view of the SFT. To support this, the SFT object pointer in the last SFT table allocated before Windows started up is declared as instance data.

The effect of a PerVMFiles=30 statement is readily seen in the following output from SFTWALK. Here, CONFIG.SYS said FILES=40, but running the SFTWALK program within a Windows DOS box showed FILES=70.

```
D:\JWD02>sfwalk
SFT @ 0116 0000 -- 5 files
SFT @ 1014 0000 -- 35 files
SFT @ 1B19 0000 -- 30 files
FILES=70
```

Thus, Windows implements PerVMFiles by linking an additional SFT table to the end of the chain, just as you saw FILES.COM do in QEMM documentation. Effectively, the DOS FILES= value has been temporarily increased without having to change CONFIG.SYS and reboot.

Windows careworn of the SFT gets even more complicated because, as noted in Chapter 1 and in the description of *Windows FileMaker* by Matt Pietrek's *Windows Internals*, KRNL386 does its own expansion of the SFT independent of PerVMFiles. When SFTWALK is run as a Windows program, see Chapter 3, the output once again changes.

```
SFT @ 0116 0000 -- 5 files
SFT @ 1014 0000 -- 35 files
SFT @ 319A 0000 -- 87 files
FILES=127
```

In this session, we happened not to execute the QEMM FILES.COM program. If we had, its added SFT would also show up in the SFT chain. As you can see, the SFT was again made to be expanded.

The source code for FILES appears in Listing 8-19. Just as SFTWALK did, FILES walks the SFTs. However, FILES does also do the SFT to get information on each open file. FILES starts with the first SFT pointer in SysArea. It opens a variable table and then goes into a loop following the SFT chain, each time adding a next table whose segment is zero or whose offset is 1 (FFFF). Only if successful, the SUM\_HANDBLES in FILES is not 0 are shown.

### Listing 8-19: FILES.C

```
/*
FILES.C List all files in DOS System File Table (SFT)
Andrew Schuelan, Revised July 1993
real mode DOS; bec files.c
protected mode Windows; bec -MS -2 -DWINDOWS files.c \undoc2\chap3\prot.c
```

This version is substantially different from that published in the first edition of UNDOCUMENTED DOS:

- Includes DOS 3.0 file structure, which is not same as 3.1
  - (LSHORT dir entry, not BYTE. Thanks to Neil Rubenking)
- DOS 2.x, 3.0, 3.1x structs combined into union, with access macros. This points up the relative inflexibility of structures struct of void, just turns into an offset, but at compile time, with no control over changing field offsets at run-time.
- Replaced previous check for possible orphaned files, which was just too flaky. Now just use simple garbage-collection test for each SFT entry, see if its owner's JFT actually contains a reference to this SFT. You can examine FILES output and decide for yourself if these files are really orphaned. If they are,

```

you can run SFT_FREE with their SFT index # on the command line.
-- Ported code to protected mode Windows (see #ifdef WINDOWS)
-- On suggestion of Geoff Chappell, got rid of AUX-CON "sanity check,"
  which was insane
-- Can be run with -FCB switch to view System FCBs instead of SFTs.
*/

```

```

#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <dos.h>
#ifdef WINDOWS
#include <windows.h>
#include <prot.h> // see \undoc2\chap3\prot.h
#endif

```

```

typedef unsigned char BYTE;
typedef unsigned short WORD;
typedef unsigned long DWORD;
typedef BYTE far *FP;

#pragma pack(1)

typedef struct file2 {
    BYTE num_handles, open_mode,
    BYTE fatr; // might be wrong?
    BYTE drive,
    BYTE filename[8], ext[3],
    WORD unknown1, unknown2,
    DWORD fsize;
    WORD date, time;
    BYTE dev_attr;
    union {
        FP dev_drv; // for CHAR dev
        WORD cluster[2]; // for disk file (BLOCK dev)
    };
    //
    // NOTE! no owner_psp!
} file2; // for DOS 2.x

```

```

typedef struct file30 {
    WORD num_handles, open_mode,
    BYTE fatr,
    WORD dev_info, // includes drive number
    FP drv, // device driver (CHAR) or BPB (BLOCK)
    WORD start_cluster, time, date,
    DWORD fsize, offset,
    WORD rel_cluster, abs_cluster, dir_sector;
    WORD dir_entry, // only difference from file31: WORD, not BYTE
    BYTE filename[8], ext[3];
    DWORD share_prev_sft,
    WORD share_net_machine, owner_psp;
    //
} file30; // for DOS 3.0 only

```

```

typedef struct file31 {
    WORD num_handles, open_mode,
    BYTE fatr,
    WORD dev_info; // includes drive number
    FP drv, // device driver (CHAR) or BPB (BLOCK)
    WORD start_cluster, time, date;
    DWORD fsize, offset,
    WORD rel_cluster, abs_cluster, dir_sector,
    BYTE dir_entry, filename[8], ext[3];
    DWORD share_prev_sft;
    WORD share_net_machine, owner_psp;
    // ...

```

```

} file31; // for DOS 3.1+

typedef union file {
    file2 f2;
    file30 f30;
    file31 f31;
} file;

// access macros
#define DOS(maj)          (_osmajor == (maj))
#define DOSVER(maj, min) (_osmajor == (maj) && _osminor == (min))
#define DOS3 FILED(p1,x) (DOSVER(3,0) ? (p1)->f30 * (p1) > f31,x)
#define DOSVER_FIELD(pf,x) (DOS(2) ? (p1)->f2 * x : DOS3_FIELD(pf,x))
#define FILENAME(pf)     DOSVER_FIELD(pf, filename)
#define EXT(pf)          DOSVER_FIELD(pf, ext)
#define FSIZ(pf)         DOSVER_FIELD(pf, fsize)
#define ATTR(pf)         DOSVER_FIELD(pf, attr)
#define NUM_HANDLES(pf) DOSVER_FIELD(pf, num_handles)
#define DEV_WORD(pf)     DOS3_FIELD(pf, dev_info)
#define DEV_ATTR(pf)     (DEV_WORD(pf) & 0x80)
#define HAR_DEV(pf)     (DEV_ATTR(pf) == 0x80)
#define DISK_FILE(pf)   (DEV_ATTR(pf) == 0)
#define DRIVE(pf)       (DOS(2) ? (p1)->f2 drive DEV_WORD(pf) & 0x003F)
#define OWNER_PSP(pf)   (DOS(2) ? 1 : DOS3_FIELD(pf, owner_psp))
#define START_CLUSTER(pf) (DOS(2) ? (p1) > f2 u.cluster[0] : \
    DOS3_FIELD(pf, start_cluster))
#define PTR(pf)         DOS3_FIELD(pf, ptr)
#define VM_ID(pf)       DOS3_FIELD(pf, share_net_machine) // Windows VM
#define DEV_DRIVER(pf)  (DOS(2) ? (p1)->f2 u.dev_drv PTR(pf))
#define DPB(pf)        PTR(pf)

typedef struct systab {
    struct systab *next;
    WORD num_files;
    file fct;
} SYS_FTAB;

void fail(char *a) { puts(a), exit(1); }

#ifdef MK_FP
#define MK_FP(seg,ofs) ((FP)((DWORD){seg} << 16) | (ofs))
#endif

#ifdef WINDOWS
#define MAP(ptr, bytes) map_req((ptr), (bytes))
#define FREE_MAP(ptr) free_mapped_linear(ptr)
#define GET_REAL(ptr) get_real_addr(ptr)
#else
#define MAP(ptr, bytes) (ptr)
#define FREE_MAP(ptr) /* */
#define GET_REAL(ptr) (ptr)
#endif

main(int argc, char *argv[])
{
    SYS_FTAB far *sys_filetab, far *next;
    file far *ff;
    char buf[9];
    int size, i, num=0, do_fcbs = 0, oeph = 0;

    // either examine SFT table or System FCB table (same format)
    int sft_ptr_ofs = 4;
    if (argc > 1 && strcmp(strupr(argv[1]), "FCB") == 0)
        ( do_fcbs++; sft_ptr_ofs = 0x1A; )

#ifdef WINDOWS
    WORD mapped;

```

```

RMODE_CALL r,
memset(&r, 0, sizeof(r)),
r.eax = 0x5200;
if (dpmi_rmode_intr(0x21, 0, 0, &r))
{
    SYS_FTAB far * far *tmp = (SYS_FTAB far * far *)
        MAP(MK_FP(r.es, (WORD) r.ebx + sft_ptr_ofs), 4),
    SYS_FTAB far *psysftab = *tmp,
    FREE_MAP(tmp);
    sys_filetab = (SYS_FTAB far *) MAP(psysftab, 0xFFFF),
}
else
    fail("Could not generate real mode 21/52!");
#e.se
_asm mov ah, 52h
_asm int 21h
_asm add bx, word ptr sft_ptr_ofs /* either SFT or System FCB chain */
_asm les bx, dword ptr es [bx]
_asm mov word ptr sys_filetab, bx
_asm mov word ptr sys_filetab+2, es
#endif

/* DOS box of 05/2 1.x doesn't provide system file tbl */
if (sys_filetab == (SYS_FTAB far *) -1)
    fail("system file table not supported");

/* could try to confirm this size by subtracting one filename
from another, as Windows does in CON CON CON CON CON code. */
size DOS(2) ^ 0x2B DOS(3) ^ 0x35 0x3B, // SFT entry size

puts("File name      Size Attr cRet Owner Cluster, DO, DPM"),
puts("-----"),

do { /* FOR EACH SFT */
    printf("\n[SFT @ %p -- %d files]\n",
        GET_REAL(sys_filetab), sys_filetab->num_files),
    /* FOR EACH ENTRY IN THIS SFT */
    if = (file far *) sys_filetab->, // DON'T MAP, already mapped
    for ( 0, i<sys_filetab->num_files, i++, num++, ((FP) ff) == size)
        if (NUM_HANDLES(ff) != 0) // don't show unused entries
        {
            printf("%-3d ", num); // SFT index
#ifdef LANDC
            // bcc pseudo printf GP faults on non-null-terminated strings
            memcpy(buf, FILENAME(ff), 8), buf[8] = 0, printf("%8s ", buf);
            memcpy(buf, EXT(ff), 3), buf[3] = 0, printf("%3s ", buf);
#e.se
            printf("%8s ", FILENAME(ff)),
            printf("%3s ", EXT(ff)),

            printf("%10u", FSIZE(ff)),
            printf("%10d", FATTR(ff)),
            printf("%d ", NUM_HANDLES(ff)),
            printf("%4s ", OWNER_PSP(ff)),

            /* A new check for orphaned files for each SFT entry,
            see if the entry is referenced in the supposed owner's
            JFT. If it isn't, it's a possible orphan. If the
            reference count (NUM_HANDLES(ff)) is 1, then it's a
            definite orphan */
            orph = 0;
            if ((_osmajor >= 3) && (NUM_HANDLES(ff) == 1) && (! do_fcbs))
            {
                WORD psp = OWNER_PSP(ff);
                if (*(WORD far *) MK_FP(bsp, 0) == 0x20C0) // real PSP
                {
                    WORD jft_size = *(WORD far *) MK_FP(bsp, 0x32);

```

```

    BYTE far *jft = *((BYTE far * far *) WK_FP(psp, Dv34)),
    int i, ok;
    for (i=0, ok=0; i<jft_size; i++)
        if (jft[i] == num) // found SFT entry in PSP's JFT
            { ok++; break; }
    if (! ok) orph++;
}

```

```
printf("orph? = DR = %d\n", orph);
```

```
if (DISK_FILE(f))
```

```
{
    printf("c %s", DRIVE(f));
    printf("25u %s", START_CLUSTER(f));
    printf("DPB %f", DPB(f));
}
```

```
else
```

```
printf("DEV %f", DEV_DRIVER(f));
printf(" %d\n", VM_ID(f)); // Windows VM # (see 2f/1683)
}
```

```
/* FOLLOW LINKED LIST */
```

```
next = sys_filetab->next, // get next.
FREE_MAP(sys_filetab), // then free old
if (next && (FP_OFF(next) != 0xffff))
    sys_filetab = (SYS_FTAB far *) MAP(next, 0xffff),
else
    break;
```

```
while (FP_SEG(sys_filetab) &&
    (FP_OFF(sys_filetab) != 0xffff)); /* ...UNTIL END */
```

```
#ifdef WINDOWS
```

```
if (!mapped & get_mapped()) != 0)
    printf("ERROR! %d mapped selectors remaining\n", mapped);
```

```
#endif
```

```
return 0;
```

```
}
```

FILES4 would be a much shorter program if there were only one SEI entry structure. Most of the complexity in this program comes from the differences between SEI entries in DOS 2.x, DOS 3.0, DOS 3.1, and higher. Most of these differences are disguised by a set of access macros so that the programmer can, for example, get the name out of a SEI entry with the simple-looking expression `FILESNAME(f)`, which behind the scenes expands into

```
(osmajor > 2) ? ((>12) filename (osmajor == 3 && osminor == 0) ?
    f->f10 filename : f->f3) filename
```

FILES4 also uses access macros to make it appear as if the SEI entry has a field containing the drive number for the open file. In fact, the SEI only contains a drive number in DOS 2.x or DOS 3.0 and higher. In SEI entries a device information word from which the `DRIVE(f)` macro extracts the drive number.

One other caveat for FILES4's length is that you can compile it for protected-mode Windows, as well as for 386-task DOS. Using `PROT_C` and `PROT_H` from Chapter 3 and a library such as Microsoft's `LIB386.LIB` (or just `LIB386.LIB`) FILES can run as a Windows program in the System VM. It's only a matter of adding `#include "map.h"` to FILES4, the `map` real-to-protected and free-mapped-to-free macros from `PROT_C` have been hidden behind yet another set of macros: `MAP_REAL_TO_FREE` and `FREE_MAP`. These macros have no effect when `WINDOWS` is not defined, that is, when compiling for 386-task DOS. This reduces the `#ifdef WINDOWS` preprocessor structures that would otherwise clog up FILES4 even worse than it is now.



However, in a new SFT is definitely an orphaned file. COMMAND.COM opened SFT to support the redirection of the TSR's output. Redirected files have two owners: see Chapter 10 for a description of how command shells provide I/O redirection. That SFT here has only one owner, COMMAND.COM, is a tip-off that the other party in the redirection hasn't existed. It's worth noting here that it's COMMAND.COM, not the TSR, that's the file's single owner, as the first edition of *Undocumented DOS* helplessly confused this point.

While the SFT contains an entry for SFT pointing back at PSP 1289h, running the FDISK program on PSP 1289h shows that this PSP's JFH has no entry pointing to this SFT entry:

```
D:\UNDOC2>hdname 1289
Files for 1289
0 ==> 1 ==> COM
1 ==> 1 ==> COM
2 ==> 1 ==> COM
3 ==> 0 ==> AUX
4 ==> 2 ==> PRN
5 ==> -1
/.
```

Note, though, that the presumed owner's JFH points to this SFT. This is how FDISK attempts to locate orphaned files. As you saw earlier, not every SFT entry has an owner which appears to be PSN's "light name" PSP, but for those that do, FDISK checks whether the owner's JFH contains a entry pointing back at this SFT entry. If the SFT entry cannot be reached from the presumed owner's JFH, and the SFT entry's reference count, NUM\_HANDELS0, is a 0, it indicates that the entry is garbage.

You can collect this garbage by specifying the helpfully orphaned SFT entry number on the command line of our next program, SFT\_FREE.C:

```
C:\UNDOC2\CHAP8>files | dos\find "OR"
3  AUX      .      0  0000      1  1289 OR DEV 0116.0048
C:\UNDOC2\CHAP8>att_free 3
C:\UNDOC2\CHAP8>files
/...
0  AUX      .      0  0000      8  9077  DEV 0070 0035
1  COM      .      0  0000      23 9077  DEV 0070 0023
2  PRN      .      0  0000      8  9077  DEV 0070 0047
```

All gone! Use this program with extreme care!

The code for SFT\_FREE.C (Listing 8-20) is amazingly simple. Once the program gets a pointer to the SFT entry it wants to free, it simply smacks a 0 into its first WORD. As seen in the structure in FDISK (Listing 8-19), the first WORD of an SFT entry is most handily setting this reference count to 0, even without changing anything else, effectively frees up the entry for reuse by DOS.

## Listing 8-20

```
/*
SFT_FREE.C -- free SFT entries specified on command line -- DANGER!!
To find SFT numbers, run FILES and look for "OR" after the owner PSP
Andrew Schulman, July 1993
*/

#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <dos.h>

typedef unsigned char BYTE;
typedef unsigned short WORD;
typedef struct _sft {
    struct _sft far *next,
```

```
WORD num;
WORD file, // SFT entries start here; num_handles is first entry
// other stuff not used here
} SFT,
```

```
void fail(const char *s) { puts(s); exit(1); }
```

```
main(int argc, char *argv[])
```

```
{
    BYTE far *sysvars,
    SFT far *sft,
    int curr_handle, handle_to_free, size, i, files, testing,

    if (argc < 2)
        fail(usage: sft free [ testing] [list of SFT handles to free] ),
    _asm mov ah, 52h
    _asm int 21h
    _asm mov word ptr sysvars+2, es
    _asm mov word ptr sysvars, bx
    if (! sysvars)
        fail("can't get SysVars"),
    #define DOS(maj) (_osmajor == (maj))
    size = DOS(2) ? 0x28 : DOS(5) ? 0x35 : 0x3B, // SFT entry size
    sft = *((SFT far * far *) &sysvars[4]),
    if (! sft) || (sft == (SFT far *) -1L)
        fail("can't get SFT");
    // walk through SFT chain once to find FILES= value
    files = 0,
    while ((FP_OFF(sft) != 0xFFFF)
    {
        files += sft->num,
        sft = sft->next,
    }
    if ((! files) || (files > 255))
        fail( "Something wrong" ),
    testing = 0,
    if (strcmp(strupr(argv[1]), "-TESTING") == 0)
        { testing++; argv++; argc--; }
    for (i=0; i<argc; i++)
    {
        if ((handle_to_free = atoi(argv[i])) > files)
            fail("Invalid SFT entry specified"),
        curr_handle = 0,
        sft = *((SFT far * far *) &sysvars[4]),
        while ((FP_OFF(sft) != 0xFFFF)
        {
            if (handle_to_free < (curr_handle + sft->num))
            { // one we want to free is in this SFT
                BYTE far *sft_entry = (BYTE far *) &sft->file;
                WORD far *num_handles;
                sft_entry += ((handle_to_free - curr_handle) * size);
                num_handles = (WORD far *) sft_entry, // first item
                if (testing)
                {
                    if (*num_handles != 0) fail("already in use!");
                    *num_handles = 1; // create phony used entry
                    // should create full-blown orphaned entry
                }
                else
                {
                    if (!*num_handles) fail("already free!");
                    *num_handles = 0; // free it
                }
            }
        }
    }
}
```



```

        // should check owner JFT to make sure really orphan!
    }
    break;
}
curr_handle += sft->num;
sft = sft->next;
}
}
return 0;
}

```

NEI\_HFILE should be necessary with the ISRs produced using the ISR skeleton in Chapter 9. As noted there, the best test for correct ISR deinstallation is the freeing up of any otherwise orphaned file handles. The generic ISR in Chapter 9 demotals with a normal DOS terminate (INT 21h) but is able to thereby eventually closing and freeing any open file handles.

### More File Handles

You've seen that the double-based file I/O routines introduced in DOS 2.0 built on two data structures, the system-wide indexed list of System File Tables and the per-process Job File Table. In contrast to the older File Control Blocks (FCBs) which applications allocated on an as-needed basis, the SFTs and Job File Tables are normally allocated by DOS itself and therefore are limited in size. You've seen that the FILEN statement in a COMBIVSN sets the number of files held in the SFTs. Normally, there are 20 possible open file handles in a process's JFT. This is dictated by the fact that a 20-byte JFT resides directly inside the PSP. This is the downside to the switch from FCBs to handles/SFTs.

DOS 3.3 introduced a function, Set Handle Count (INT 21h) AH=67h, which can increase the size of the calling process's JFT, thereby increasing the files and devices that may simultaneously be opened using double-based file I/O. Some serious programmers claim the function "doesn't work" merely because they forgot to increase the FILEN setting when attempting to keep 50 files open at once, or because they think that calling this function automatically allows them to open 500 files with a single file handle for each such as open() (see FILEN vs File Handles' earlier in this chapter). In either case, these supposed bugs are nothing more than concept errors.

However, as indicated by the applicable entry for INT 21h AH=67h, this function "sets only in this function if the user calls it on AH=67h." *Tech Journal* article (April 1988) also stated that the function can be used to "force 64k free memory because its code uses an RCR instruction instead of the correct RCR."

It remains to be just as easy to perform the same function course I. Since a JFT is currently held directly at offset 18h in the PSP, it seems likely it should be difficult to increase its size. However, since DOS 3.0, the PSP has also contained a file pointer to the JFT and to a word holding its size. The relevant fields in the PSP, which we already used in the orphan seeking code in FILEN.C, are:

```

18h  20 BYTES  DOS 2+ JFT
32h  WORD    DOS 3+ max open files
34h  WORD    DOS 3+ JFT address

```

You can't do anything to increase the size of the array at offset 18h in the PSP, but you can allocate a new target stack area or copy for the JFT, bump up the count at offset 32h in the PSP, and table it to the new one, and then set the pointer at offset 34h to the new table. The same type of manipulations you've become familiar with seemingly static DOS arrays, such as the SFTs and CDBs, whose far pointers are located in SYSVAR.

The program in Listing 8.7, FILEHANDLE.C, carries out this series of operations. FILEHANDLE takes a number on its command line and attempts to resize its own JFT.

```

D:\JNDOC2>handle 40
Currently 20 max JFF file handles
And SFT FILES=40
Max file handles increased to 40
Opened 34 files
    
```

To test that more files can be opened, FHANDLE continually opens its own executable file arg[0], in a loop until the Microsoft's dos.open() function fails.

### Listing 8.21: FHANDLE.C

```

/*
FHANDLE.C
Alternative to using INT 21h function 67h (added in DOS 3.3)
Andrew Schulman, revised July 1993
*/
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <dos.h>

typedef unsigned char BYTE;
typedef unsigned WORD;
typedef unsigned long DWORD;
typedef BYTE far *FP;

#define MK_FP \
#define MK_FP(seg,ofs) ((FP)((DWORD)(seg) << 16) | (ofs))
#endif

extern unsigned files(void); // In fh.c
void fail(char *s) { puts(s); exit(1) }

typedef struct _sft {
    struct _sft far *next,
    WORD num,
    // other stuff not used here
} SFT;

WORD files(void)
{
    SFT far * far *sysvars; // treat SysVars as array of far ptrs
    
```

- NOTES
- Generally need FILES greater than 20 in CONFIG.SYS to have any effect. Actually, this isn't strictly necessary. Certainly you could have a larger JFF than your SFT and have multiple JFF entries pointing to the same SFT entry. This occurs already, where normally file handles 0, 1, and 2 all correspond to SFT entry 0 (CON). But this generally isn't why programmers want to increase their JFF size! So, if your intended JFF size is larger than the current SFT size, you can:
    - Change FILES in CONFIG.SYS and reboot.
    - Or use Quarterdeck's FILES.COM program to increase SFT size.
    - Or see XLASIBRV program in Chapter 2, which expands CBS, and adopt it to expand the SFT chain.
  - If you want to open >20 files simultaneously with high level function like fopen(), you probably must increase run-time library tables, for example, in Microsoft C the size of table for FILE\* is hard-wired to 20 in %s\source\startup\\_file.c (#define \_NFILE\_ 20). So increase the size of this table as well (and recompile startup.c) if you want to use >20 fopen() at once, even after using the following code.

```

/*
#include <std. h>
#include <stdio.h>
#include <string.h>
#include <dos.h>

typedef unsigned char BYTE;
typedef unsigned WORD;
typedef unsigned long DWORD;
typedef BYTE far *FP;

#define MK_FP \
#define MK_FP(seg,ofs) ((FP)((DWORD)(seg) << 16) | (ofs))
#endif

extern unsigned files(void); // In fh.c
void fail(char *s) { puts(s); exit(1) }

typedef struct _sft {
    struct _sft far *next,
    WORD num,
    // other stuff not used here
} SFT;

WORD files(void)
{
    SFT far * far *sysvars; // treat SysVars as array of far ptrs
    
```

```

SFT far *sft;
WORD files = 0;
_asm mov ah, 52h
_asm int 21h
_asm mov word ptr sysvars+2, es
_asm mov word ptr sysvars, bx

if (! sysvars) return 0;
for (sft = sysvars[1], FP_OFF(sft) != 0xFFFF, sft = sft->next)
    files += sft->num;
return files;
)

main(int argc, char *argv[])
{
    WORD far *pmax = (WORD far *) HK_FP(_psp, 0x32),
    BYTE far * pjit = (BYTE far * far *) HK_FP(_psp, 0x34),
    BYTE *new_jft,
    WORD max = *pmax,
    WORD new_max,
    int f, i;

    if (argc == 2)
        fail("usage: handle [-testing] [#handles]"),
        new_max = atoi(argv[1]),
        printf("Currently %u max JFT file handles\n", max),
        printf("And SFT FILES=%u\n", files());

    if (new_max <= max)
        fail("nothing to do");
    else if (new_max > files())
        fail("FILES= too low edit CONFIG SYS and reboot\n"
            "or run a program like BERM FILES.COM to grow the SFT"),

    if (! (new_jft = (BYTE *) malloc(new_max)))
        fail("insufficient memory");
    _memcpy(new_jft, pjit, max), // copy over old entries
    _memset(new_jft+max, 0xFF, new_max - max), // fill in new entries
    *pmax = new_max; // set new max file handles
    *pjit = (BYTE far *) new_jft; // set new JFT

    printf("Max file handles increased to %u\n", new_max),
    // now test how many files we can open by opening ourselves (argv[0])
    for (i=0; i++)
        if (!_dos_open(argv[0], 0, &i) == 0)
            break,
    printf("Opened %d files\n", --i);

    if (argc > 2 && strcmpi(argv[2], "-TESTING") == 0)
    {
        BYTE cmd[16],
        // printf("Closing %d => %d\n", f, new_jft[f]),
        _dos_close(f), // close last one so we can spawn shell
        sprintf(cmd, "H2NAME .04X", _psp),
        printf("\n> %s\n", cmd); // show command line
        system(cmd); // run H2NAME on my PSP
        printf("\n> FILES\n"),
        system("FILES"); // run FILES
    }

    return 0;
}

```

If HANDLE is also run with the TESTING command line option it runs the previous H2NAME (Listing 8-18) and FILES (Listing 8-19) programs first it closes the most recently opened file set of files, it then forks a child from which to spawn the programs. The only HANDLE and

sometimes has problems spawning H2NAME and FILES. This not only helps verify that you can really open more files, but also helps to show how these three different programs tie together.

```
D:\UNDOC2>fhandle 40 -testing
Currently 20 max JFT file handles
And SFT FILES 40
Max file handles increased to 40
Opened 34 files
```

```
> H2NAME 6C7B
Files for 6C7B
0 ==> 1 ==> CON
1 ==> 1 ==> CON
2 ==> 1 ==> CON
3 ==> 0 ==> AUX
4 ==> 2 ==> PRN
5 ==> 3 ==> FHANDLE EXE
6 ==> 4 ==> FHANDLE EXE
...
37 ==> 35 ==> FHANDLE EXE
38 ==> 36 ==> FHANDLE EXE
39 ==> -1
```

```
> FILES
# Filename Size Attr cRef Owner Cluster, BB, DPB
[ SFT @ 0116 00CC -- 5 files ]
0 AUX 0 0000 10 9077 BEV 0070 0035
1 CON 0 0000 29 9077 BEV 0070 0023
2 PRN 0 0000 10 9077 BEV 0070 0047
3 FHANDLE EXE 9126 0020 1 6C7B 0 30714 DPB 0FF1 0000
4 FHANDLE EXE 9126 0020 1 6C7B 0 30714 DPB 0FF1 0000
[ SFT @ 1014:0000 -- 35 files ]
5 FHANDLE EXE 9126 0020 1 6C7B 0 30714 DPB 0FF1 0000
6 FHANDLE EXE 9126 0020 1 6C7B 0 30714 DPB 0FF1 0000
... etc. ...
36 FHANDLE EXE 9126 0020 1 6C7B 0 30714 DPB 0FF1 0000
```

You might wonder why there are all these SE entries for the same file. From the FILES output, they all look the same. However, one of the fields in the SEI not shown by FILES is the current file position. This can not be shared among multiple opens of the same file, since the following code (for example) is valid:

```
int f1 = open("foo.bar", O_RDONLY);
int f2 = open("foo.bar", O_RDONLY);
lseek(f1, 10, SEEK_CUR);
lseek(f2, 20, SEEK_CUR);
```

The f1 and f2 handles must lead to separate SEI entries. Even loading SHARE can't change this, assuming the second `open()` is allowed to succeed.

Once again, increasing the DOS file handles does not automatically allow you to open more files with the `open()` function. If you need to increase the number of openable files, consult your compiler's startup code. For example, in Microsoft C you can increase the FILE\* maximum by changing the value of `NEHE` in the startup code (see `STARTUP.PCR10D.VI.ASM`).

Note that the original FILE is not inheritable, so `HANDLE` can't pass its increased wealth along to any child that might be created. For an in-depth look at this topic, see the article, "DOS File Handle Limits" by David Bark, *THE SYSTEMS* (February 1993; this magazine is now the *Windows/DOS Developer's Journal*).

## System FCBs

And now for a trip down memory lane.

So far we have seen how handle-based DOS file operations manipulate SEFs. But remember DOS File Control Blocks introduced from CP/M when it was the standard of the 8-bit microcomputer world. (That's why DOS handles as it were, FCBs.) Recall keeping some nice, medium-nearly-1000-FCB work. Instead of getting back a handle, which causes the operating system to open a file, an application creates an FCB *in memory* (data area) and tells the operating system where this FCB is. DOS refers to the FCB in the application's data area.

Thus, the application allocates FCBs on behalf of the operating system, rather than the other way around. Soak—the problems with keeping crucial operating system data in the application's data area should be rather obvious. For example, whereas today's DOS can close any open files when an application exits in the face of FCBs, the operation system couldn't do anything about programs that left behind open files when exiting.

Furthermore, putting the FCB under the application's control meant that every application using FCBs was dependent on its precise structure. Contrast the way that (for example) the vast majority of DOS applications are not dependent on the structure of an SEF. The dependence of applications on the FCB structure has obviously made it impossible to change this structure, even when necessary. (In instances, to support media larger than 32 megabytes, FCBs in fact are a perfect example of *vertical* compatibility—expose operating system internals to applications.)

What is less obvious is that there was one benefit to FCBs over the new, improved handle-based file functions. Since the application allocates FCBs, there are practically unlimited open files. An FCB system doesn't need a HFN—setting an OPEN SYN. When the application needs to open a file, it creates another FCB.

SEFs and FC handles were introduced in DOS 2.0, and with them the need for a HFN—setting so that users could exercise some control over the trade-off between having lots of potential open files on the one hand and taking too much precious memory on the other.

However, Microsoft could not get rid of FCBs entirely. While over time Microsoft has reduced its support for FCBs—for example, in the Windows DOS extenders, the MS-DOS programmers reference continues to document the FCBs, while advising that as superseded. According to Microsoft, programmers "should not use a superseded function except to maintain compatibility with versions of MS-DOS earlier than version 2.0."

Nevertheless, Microsoft use continues to use FCBs, even in parts of DOS 5.0 and 6.0 that (presumably) do not require compatibility with DOS 1.0. Microsoft seldom follows its own advice: this advice is intended, apparently, for everyone else. For example, a simple INTRN.PY script reveals that COMMAN.COM extensively uses FCB functions. It uses `format` on 29th FCB, `ParseFileTime` as part of its normal operation, `constrains` (1th, 10th, and 12th FCBs). Next, for the DIR command (perhaps a clever way to get starting cluster values), `function` (8th FCB), `Dir` (6th FCB), and `function` (7th FCB). Reason, for RFN with wildcards, `function` extended FCBs are necessary to change `constrains` to support DOS 4.0 disks (disks formatted in order DOS 4.0 or higher can have their volume labels changed to special names). Media ID generic, `IOCTL` calls `constrains` earlier.

So now, how DOS could use to support FCBs without totally compromising the system. By providing a special System FCB. When a program uses FCBs, DOS copies all pertinent information from the program's own FCB into an actual System FCB, does the actual work using the System FCB, and then copies the information back into the user's original FCB before returning control. Programs using FCBs do not have to know anything of System FCBs.

Since FCBs refer to the application System FCBs, each has a `own` or `id`. However, System FCBs use SEF entries to do so. The limit for System FCBs and SEFs is identical. Thus, at the interface, DOS can use to work with SEF entries work exactly the same way as with SEF FCBs. The needed difference are that System FCBs are kept on a separate list from the SEFs, at a clear

pointer to the System FCB chain is at a different location in SysVars from the SFT pointer, and that DOS must keep a System FCB synchronized with the old FCB that its owner believes is controlling the FCB. Since FILEN determines the number of SFT entries, so FCBs determines the number of System FCBs.

That System FCBs are SFTs in disguise makes it easy to use our SFT-walking programs to view System FCBs. If the standard DOS FCBs is enabled, `MSIWAJ Listing 8.17` can determine the value of FCBs by walking the chain of System FCBs, just as `FILEN` determines the value of FILEN by walking the SFT chain.

```
D:\>DIRC2>SFTWALK
SFT @ 0116 00CC -- 5 files
SFT @ 0114 00DD -- 35 files
FILES=40
System FCB table @ 1979:0000 -- 8 files
FCBS=8
```

The only difference of the SFT-walking code is that SysVars[4]h rather than SysVars[4] is converted to SFT.

```
printf( "FILES=%d\n", sftue.k16, SFT );
printf( "FCBS=%d\n", sftue.k1021A, System-FCB table );
```

Note that `FILEN` Listing 8.19 takes an optional FCB switch in its command line.

The problem with System FCBs, however, is that because the FCBs statement in `CONFIG.SYS` has the effect of making System FCBs programs or programs upon unlinked FCBs. Once the System FCBs are unlinked, DOS says to use a most-recently used System FCB. There was a complicated procedure for that in IBM's DOS 4.0 that involved protected FCBs (the `FCBS=SYS` statement) that was dropped in DOS 5.0.

While this may include a lot of extra work just to support a archaic manner of opening files, it does permit programs that use a older calls to coexist with the newer techniques. Hard to believe that people can be many years still use FCBs for one reason or another. System FCBs give DOS the capability to do file sharing, networking, and other purposes while still allowing ancient programs to run without change.

One reason to be careful about System FCBs is that a heavily scripted piece of code in Windows 4.1.4 would use VARD code to see the offset of the first System FCB header as part of its test for generic MS-DOS. Any version of DOS in which this offset is non-zero (that is, `!P OFF SysVars[0x1A] != 0`) Microsoft considered a competitor's DOS, and certain versions of Windows issue a warning error message. Chapter 1 discusses this in great detail (see MSDN FCB in Chapter 1).

## The SHARE Hooks

In our final discussion of the DOS file systems you may have seen something where SHARE fits in. How does it fit? SHARE patches itself into the DOS file system.

If you go through the appendix to this book, you will see many, many references to SHARE. Far more than you would expect from what seems like a fairly simple file sharing and redirection program. While the network redirection interface that we describe next is no mode of software engineering compared to SHARE, it seems downright elegant.

While SHARE provides a small set of API functions, some documented (using INI 21h AH=51h and INI 21h AH=10h) this is not its primary mode of operation. Instead, SHARE patches itself into the DOS kernel.

The DOS kernel makes calls to the far pointers at various significant times during execution. These pointers, known as the SHARE hooks, are located at negative offsets from the first SFT. By default, the SHARE hooks point at one of two dummy routines, which merely set or clear the carry flag, depending on whether the routine should succeed or fail when SHARE is not loaded. SHARE

changes these pointers to transfer control to its own routines. As a result of SHARE's modifications to the DOS kernel, some INRTRM000 programs can't remove SHARE. DOS will hang the next time it tries to call one of the SHARE functions through the no longer valid pointers.

SHAREHOOK.C (Listing 9-22) is a simple program to dump out the SHARE hooks. Usually mere inspection of the SHAREHOOK output tells you whether SHARE is installed because prior to loading SHARE all the hooks point to the same two routines; after loading SHARE each hook points to a unique location, which can of course be disassembled.

```
C:\VUNDOC2\CHAP8>sharehook
-3c 0116 0090 FDCB:44AA unknown
-38 0116 0094 FDCB:44AE OpenFile
-34 0116 0098 FDCB:44AE CloseFile
;
-c 0116 00CD FDCB:44AE CloseFileIfDup
-8 0116 00C4 FDCB:44AA Close*
-4 0116 00C8 FDCB:44AA UpdateDirIn$FT

C:\VUNDOC2\CHAP8>share
SHARE installed
C:\VUNDOC2\CHAP8>sharehook
-3c 0116 0090 D000:0000 unknown
-38 0116 0094 1842:0954 OpenFile
-34 0116 0098 1842:0958 CloseFile
;
-c 0116:00C0 1842:0980 CloseFileIfDup
-8 0116:00C4 1842:0984 Close*
-4 0116 00C8 1842:0988 UpdateDirIn$FT
```

### Listing 8-22: SHAREHOOK.C

```
/*
SHAREHOOK.C -- Dump DOS SHARE hooks
Andrew Schulman, June 1993
*/
#include <stdlib.h>
#include <stdio.h>
#include <dos.h>
typedef void (*FJNCPTR)(void);
char *share_func_name[] = { "unknown", "OpenFile", "CloseFile",
"CloseAllMachine", "CloseAllProcess", "CloseByName", "LockRegion",
"UnlockRegion", "CheckRgnLocked", "GetOpenFileListEntry",
"update FCB from $FT", "get first cluster of FCB?", "CloseFileIfDup",
"Close*", "UpdateDirIn$FT" };
main()
{
    unsigned char far *sftptr,
    FJNCPTR far *sharehooks, far *hook;
    int ofs, i;
    _asm mov ah, 52h
    _asm int 21h
    _asm les bx, dword ptr es:[bx+6]
    _asm mov word ptr sftptr+2, es
    _asm mov word ptr sftptr, bx
    sharehooks = (FJNCPTR far *) (sftptr - 0x3c);
    for (i=0, hook=sharehooks, ofs=0x3c, i<15, i++, hook++, ofs+=6)
        pr printf( "4x 2Fp ?Fp 1s\n",
        ofs, hook, *hook, share_func_name[i]);
    return 0;
}
```

The three most important SHARE hooks are those called when a file is opened, cloned, and closed. Additionally, hooks are used to pass through several undocumented DOS calls for cloning files by process, to lock or to unlock regions of a file for checking whether a region is locked. As an example, hooks where use is not entirely known. A list of the SHARE hooks are listed in the appendix under INT 21h function 52h.

As noted earlier, to open a file, DOS first canonicalizes the filename. After determining that there are no clones, it searches the caller's file table for an open SEI entry. DOS then calls the SHARE file open hook. In any DOS calls through whatever function pointer is plugged in at offset 36h from the start of the file's SEI. If the file name matches with carry set, the file is not available due to sharing restrictions, so DOS deems and tries calling the hook until the sharing retry count is exhausted or the file becomes available. SHARE proceeds by beginning a DOS critical section, checking whether the file may be opened. Calling in the SHARE records of the SEI entry it secondarily ending the critical section, so determine whether the file may be opened. SHARE searches its list of open files for a match with the name of the file. If one is found, the file open modes of the new open of the file in the context of original operations compared according to the standard SHARE rules; for example, an open (O) Write (W) beats a subsequent open (O) Read (R) Write (W). If it matches is found, a new sharing record is created for the new file and the open is always allowed to succeed. When doing this SEI holds, SHARE stores, among other things, the file's owner and links the SEI to any SEI already referencing the same file.

When a file is closed, DOS again calls on SHARE. After beginning a DOS critical section, SHARE removes any locks placed on the file by the calling process and removes the SEI entry from the chain of SEI entries for the file, if the close removes the last reference to the SEI entry. Finally, SHARE calls the DOS critical section (DOS SEI) to destroy the reference to the file. SHARE also erases the sharing record for that file.

DOS also calls on the three most important SHARE hooks when the file size or time stamp changes. DOS calls on SHARE to propagate these changes, as well as any change to the starting cluster number. A call to the SEI structure at the 500h offset. This causes that all processes have a consistent view of the file, especially the sharing pointer and device. You can see an interesting result of this hooking when cloning a file under a debugger. From another window of the debugger, I can copy a file and see the pointer in the debug window not even though its directory entry still shows a zero cluster provided that the second program reads until DOS returns a valid file handle. I can seek to the end of the file, either through the seek or the directory entry.

It is important to use SHARE correctly during a pre-emptive multitasking system such as Windows. In order to use the SHARE to intercept file requests from multiple Windows VMs, you must run it early. Windows, being SHARE, runs within a Windows DOS box, a bad idea, so Microsoft prevents you from doing this.

Nevertheless, Microsoft prevents users from running SHARE in a Windows DOS box by preventing the SHARE to be early installed. Whether or not SHARE is loaded before Windows, the SHARE detect (INT 21h AX=4000) always succeeds after Windows. This unfortunately has led to the awkward situation that Windows automatically loads SHARE or that Windows contains built-in support for SHARE. After Windows for Workgroups 3.11 and in Chicago, Windows does contain built-in support for the SHARE hook of the VSHARE 386 VxD. In fact, Windows just works to succeed to run the SHARE in a DOS box. By hooking the INT 21h SHARE detect, the call, Windows intercepts the DOS MGR VxD call into WIN 386 INT. It takes SHARE INT into not installing SHARE. The call sees that INT 21h call sees that SHARE is supposedly already installed and calls DOS MGR sees the same INT 21h hook installed with the VMM function Hook V86 Int. Check to see if you are using FASTOPEN and NLSFUNC.

You can run SHAREHOOK to check Windows to confirm whether SHARE is really installed or



```

C:\WINDOWS\SYSTEM32>share
SHARE already installed
C:\WINDOWS\SYSTEM32>share -hook
-3c 0116 0090 FDCB:44AA unknown
-38 0116 0094 FDCB:44AE OpenFile
-34 0116 0098 FDCB:44AE CloseFile
-30 0116 009C FDCB:44AA CloseAllMachine
-2c 0116 00A0 FDCB:44AA CloseAllProcess
-28 0116 00A4 FDCB:44AA CloseByName
? -

```

That table reveals different hook functions are installed as a sure tip off that SHARE is not already installed. Again, none of this applies under VSHARE 386.

Microsoft's DOTS programmer's reference explicitly mentions the fake SHARE problem: "Some operating environments, such as Windows, intercept this multiplex interrupt and always return a nonzero value, but the Share program is loaded or not." It is sick when, due to the manufacturer's own actions, one can use a function labeled Get SHARE FN Installed State to get the SHARE installed state.

However, Microsoft does recommend a true SHARE-detection method: "To determine whether file sharing is available, a program should check for error values upon returning from carrying out a file sharing function, such as Lock, UnlockFile, Interrupt 21h (fnct on 5C h)." It would be nice if PC debug tools or programs would heed this advice, rather than blindly reporting "SHARE installed" under Windows. The code in Listing 8-23, IS SHARE.C, presents both the non-bogus SHARE-detection algorithm that actually works based on function 5C h.

### Listing 8-23: IS SHARE.C

```

/*
IS_SHARE.C -- Determine if SHARE really loaded
Andrew Schulman, April 1993
*/

#include <stdlib.h>
#include <stdio.h>
#include <dos.h>

int bogus_share(void)
{
    unsigned char bogus;
    __asm mov ax, 1000h
    __asm int 2fh
    __asm mov bogus, al
    return (bogus == 0xFF);
}

int is_share(void)
{
    __asm push si
    __asm push di
    __asm xor bx, bx
    __asm mov cx, bx
    __asm mov dx, bx
    __asm mov si, bx
    __asm mov di, bx
    __asm mov ax, 5C00h /* try to lock */
    __asm int 2fh
    __asm jc lock_error
have_share:
    __asm mov ax, 5C01h /* SHARE there: unlock! */
    __asm int 2fh
    __asm pop di
    __asm pop si
    return 1;
}

```

```

lock_error:
    asm pop di
    _asm pop si
    asm cmp ax, 1 /* Invalid function: no SHARE */
    _asm ne have_share_2-
no_share
    return 0,
Have share ?
    return 1,
}

main(int argc, char *argv[])
{
    int share = ss_share(,,
    if (bugas share(),
        puts( share ? "SHARE loaded"
                "SHARE supposedly loaded, but it really isn't" ),
    else
        puts( share * "SHARE supposedly not loaded, but Func 5Ch present" :
                "SHARE not loaded" ),
    return 0;
}

```

In a Windows DOS box, unless SHARE was run before starting Windows, IS SHARE displays the message "SHARE supposedly loaded, but it really isn't."

Microsoft's VSHARE 386 makes this problem a thing of the past. As part of its major Chicago project, DOS 3.0 and Windows 4.0, Microsoft is introducing many, possibly all, pieces of MS-DOS protected by SHARE into Windows 386. In a protected mode, set of device drivers (see Chapters 1 and 3). Microsoft's project does some 400 new VxDs, including VSHARE 386, as part of Windows for Workgroups 3.11 prior to the release of Chicago. VSHARE 386 provides file sharing and locking, without having to run SHARE (still in Windows). Because it's a VxD, VSHARE also does not waste conventional real-mode A: disk 32-bit protected mode code allocated out of extended memory. With a DEVICE=VSHARE 386 statement in SYSTEM.INI when Windows reports that SHARE services are available, they actually will be. What a concept!

## The MS-DOS Network Redirector

As we saw, MS-DOS contains a set of function pointer hooks that are set by SHARE 1.1E and that MS-DOS calls to customarily point during the I/O. Another set of hooks in MS-DOS is the network redirector interface. When MS-DOS calls SHARE 1.1E, it requests function pointers with a far CALL. DOS calls network redirector's INT 21h calls. It is, where, is SHARE requests setting the value of time, then points to the DOS kernel, where a redirector requests (looking INT 21h).

Normally, to configure Event I/O redirector, see Chapter 10, the MS-DOS network redirector is a mechanism for connecting a computer into the stream of file system and pointer requests so that some requests can be sent out onto another computer, a special way for sending a redirector services such requests by using the network packets sent to a file server, but there is nothing in the redirector mechanism that restricts it to networking. In this section, we are concerned with the file system, rather than pointer, aspects of redirection.

Microsoft uses the redirector interface since DOS 3.11 to allow DOS programs to transparently access all of the systems. This was completed with the version of DOS to first support networking is no coincidence, since the redirector interface is the vehicle DOS provided for the implementation of network services such as IBM PC LAN.

Microsoft uses the redirector interface to provide CD-ROM access under DOS through the Microsoft CD-ROM File System, MMSDFX. Other services of the interface include network operating systems like Banyan and 3Com NetWare 1 to uses the redirector and, starting with version 4.0, so does Novell NetWare.

It is worth stressing the usefulness and power of the redirector interface. Although these inconsistencies do exist in places, as you will see, the possibilities for its application are enormous. In fact, Microsoft has not documented it and will not support it. In response to one customer's request, that director information is posted on its Online technical support service. Microsoft responded:

The FN\_200 interface to the network is an undocumented interface. Only FN\_200 (the redirector DLL) is in the installed state of the network services as documented.

So, a third-party vendor reverse-engineered and documented the interface to "Linking in local DOS" to SysShare via Ability in Windows, but Microsoft provides almost no help or support for programming on that API, and we do not guarantee that the API will exist in future versions of MS-DOS.

The party line is, apparently, "Here's where you get the info, but you better not use it."

Reason I say that the redirector is not even documented is when Microsoft, except for a handful of letters (the "b" that had Microsoft's logo) made redirector source code and supporting materials selectively available, thus making the network redirector a good example of how supposedly undocumented functionality is available, source code downloaded, but not fully documented. Such discriminatory documentation is far worse than any outright complete absence of documentation.

## Novell's Undocumented Redirector Interfaces

**So that we're not dumping unfairly** on Microsoft, it's worth pointing out that its major competitor, Novell, has its own undocumented redirector interface (two of them, in fact). NetWare 3.x and 4.x has an undocumented server interface for Name Spaces (Macintosh HPFS, etc.). The name spaces are supposed to be dynamically allocated, but Novell hard codes the hook in slots. Name Space NLMs (NetWare loadable modules, sort of a cross between DLLs and VxDs) then plug in to these pre-allocated slots.

Novell's NLM SDK describes a second redirector-like interface: "A new feature in NetWare 4.0 is the volume switch, which makes it possible for NetWare to access non-NetWare file systems, such as the NFS file system." While Novell mentioned this Volume Switch at one of its BrainShare conferences, it declined to release developer documentation. According to one engineer, they [Novell] like to handle this kind of thing strategically, which means they bring in the developer who is whining about it and give them documentation under strict NDA."

In a sense, the information and techniques presented in this chapter were, originally, derived from my personal, considerable experience with the redirector since the first edition of *Undocumented DOS*.

Because of the possibilities for confusion in terminology, we need to define some terms. The term *redirector* (or *file redirector*) means the functionality and hooks DOS provides for foreign file systems and file systems that use DOS's own programs that "speak" interface. Most, strictly, DOS interface terms call it the FN\_200 hook or FH; this is called the redirector interface. Programs can receive the use of DOS's own interface, FN\_200 hook or FH, such programs are called *redirectors*.

After the *BrainShare* conference, I was struck by how available the term "redirector" is. Since then, I've found that the term "redirector" may appear widely. IBM introduced IFS as part of OS/2. IFS allowed a user to have file systems to be developed and to run seamlessly, not by operating system files. IFS is similar to the OS/2 IFS (High Performance File System). IFS is used as a replacement for the old EFS file system, another OS/2 IFS user was the Novell NetWare Redirector.

In essence, IFS was to be a legitimized version (under OS 2) of the hidden, undocumented redirector interface (under DOS). The fact that the original IFS appeared frequently in DOS 4.0 documentation indicates that IBM's intent was to provide a stepping stone to help developers migrate toward OS/2 (particularly LAN Manager).

The IFS interface appears to have been implemented under DOS 4.0 using the IFSEFS.C program, which made FileSrv a redirector. In other words, IFS was an extra layer on top of the redirector. Some additional subfunctions (not appeared under the redirector interface) to support what may be a host of structural rework (if above) the existing redirector interface.

However, the reality is that IFS in the DOS world received no publicity at all, and DOS 4.0 was never really back to the mixer. There is not much benefit to a further discussion of IFSEFS.C. REDIRFS.C (or any other DOS 4.0 specific specific on the network redirector). We consider IFS here only insofar as it's still usable and overlaps the redirector interface. In other words, we'll refer to subfunctions of IFS 2H function 1 (redirector interface subfunctions) even if they were first introduced in DOS 4.0 as IFS subfunctions.

In Chicago, Microsoft will introduce an IFS Manager, IFSMGR.386 and IFSEFIP.SYS. This, rather than its network redirector, will be the approved way for writing installable file systems. According to Microsoft, the benefits of IFSMGR over the redirector are: 1) it will be documented properly (in the Group 13.0K 2) it clearly handles the case of multiple redirectors, and 3) it serves as a common installation point with DOS global variables at the SYM and so on.

In the remainder of this chapter, we delve further into what happens inside DOS file I/O calls using INTRP.Y to see the various instances under such DOS calls a redirector. A detailed specification for the OS 2H API 11h interface is presented, along with an extensive sample program (the Phantom). Readers of the first edition of *Undocumented DOS* should note that the rewritten PHANTOM.C incorporates far more extensive knowledge of the redirector interface than did the original PHANTOM.PAS. PHANTOM.C is two large source files (see notes in context), but we present a detailed disk spec. The source code for PHANTOM files selected file I/O calls such as Open, Read, Write, and Mkdir. Complete source code is presented on the accompanying disk in 3.13.DOC 2C.HAPPHANTOM.C.

Several files of *Undocumented DOS* provided important additional information and clarification on the redirector interface. Thanks go due to Carsten Bokrosch, Andersen, Dave Andrews, Dwight Bales, Eric Bergste, Eric Lark, Mike Kates, David Markan, Kurt Pate, Mitchell Scholstein, Mike Nuech, and Martin Westmeier.

For additional information on the redirector interface, readers may wish to read the article, "A DOS Redirector to MS CD-ROM" by Tom Harper, *The Linux Journal*, March 1993. Another interesting source of information on the redirector interface from the DOS emulation code (MIDOS in Cary's McBot University's Mch operating system) is the relevant file is DOS.FSC, available by ftp from <http://csrc.nyu.edu>.

## Using the Network Redirector Interface

In essence, we want to use the redirector interface to manufacture DOS drives. A DOS drive is any entity that has a drive letter in a C:\PS entry, and that behaves like a hard drive or floppy, in that you can access it with normal DOS disk I/O commands and file commands and INE 21h function calls. Whether there is magnetic media or the other kind is irrelevant, since a file system is a purely logical construct.

One can view a (almost any) area of computing as a set of file operations. To take one example, say that every evening you log on to an information service, such as CompuServe, with a hierarchical structure of forums, message and library areas, subforums, topics within message areas, and so on. Normally, to access such services, you would dial up through the modem and send the service various command strings or menu selections. However, it might be more convenient to pretend that the information service is just another drive on your machine. Changing to drive J and typing CD

IBM DOS SECURITYS, for example, might be equivalent to joining the IBM DOS SECURITYS device on BIX. Conceptually, it is quite simple. All you need is to designate a drive letter for the information service, install a piece of software that catches all disk, directory, and file requests sent to that drive, and have the software service these requests by issuing the proper Cmpfsrv or BIX mechanisms. With a sufficiently fast modern PC, this would feel just like using a disk.

**Front-End Hooks and Device Drivers vs. Back-End Redirectors** There are several ways to catch INE 21h file I/O requests. One method does quite literally "catch" disk, directory, and file requests by hooking INE 21h and watching every function call that comes in. Novell NetWare earlier than 4.0 works this way (see NETV in Chapter 4). Because this method sees INE 21h calls before DOS has a chance to DOS read or see some of them, it all... it is sometimes called a *front-end hook*. Novell NetWare 4.0 and 4.11, because NETV puts a screen around the DOS APIs.

A front-end hook is conceptually simple (see FILEC00132.C and DOSVIRG in Chapter 2), but there are problems with hooking INE 21h and looking for all file- and directory-related calls for your speciality. For one thing, you must deal separately with ECB and handle-based calls. DOS file functions available to applications provide two very different interfaces. As you saw earlier in this chapter, the older ECB calls have the following implications:

- **pathnames**—In a less convenient open files because the state for the open file is kept in a user-supplied structure.
- **wildcard file name specification** in Delete and Rename.
- **FCB functions cannot use pathnames**.

The file handle calls imply that

- **the concurrent open file count** is limited to the FILES= limit in CONFIG.SYS, that is, to the size of the SFT—and for any given application, to the size of its FCB.
- **filenames can contain pathnames**.
- **simplicity of use**—image cookies. Fileless applications don't have to know the underlying FCB and SFT structure.

The redirector interface unifies the two access methods so that, with one exception (discussed later), a redirector need not know by what method a file is being accessed. This shows the interface functioning in a way very different from what it is as a back-end to DOS. This access method independence is a great advantage, but it runs the danger of the redirector interface overloading INE 21h replacement as the means of representing different file systems. Rather than trying to duplicate the entire range of file system calls in the DOS programming interface, including both ECB and handle-based file access methods, a redirector instead plugs in at a level where a much of the higher-level administration calls to the file calls become transparent. Underlies with a homogeneous interface. For example, all file name strings used to open and otherwise manipulate files arrive at the redirector as fully resolved paths; the DOS kernel has already done the work of resolving the drive and directory.

One major advantage with hooking INE 21h is that you must implement your own version of function 4h (FIC). While Microsoft does provide a Set Execution State function (INE 21h AX 4B05h) that allows you to replace the current FIC, there is little additional support. Each new version of DOS requires you to change to one EXEC code. In fact, it is reported that one reason for Novell's strategy of hooking INE 21h to redirectors is INE 21h is that the company is planning on the DOS 5 changes to EXEC.

Another problem is that a front-end hook requires a fairly large amount of state, just to correctly handle complex path names.

Finally, when you use the redirector support for undocumented calls such as `Filename INT 21h AH:00h`, support appears automatically. A front-end hook must explicitly code in support for all undocumented calls. With a redirector interface, DOS takes care of cooking all file system requests for you.

On the other hand, the redirector interface does not support undocumented, so it may not seem to hold any extreme advantage over hooking INT 21h. And of course, getting total support for DOS calls, both documented and undocumented, requires implementing the entire redirector interface. For example, support for INT 21h AX:020h Get Assign List Entry, see `NTDRVC` in Listing 8.15, requires that the redirector support INT 21h AX:111h Dr Reduction. In any case, the point remains that the redirector provides a more homogenized, narrower interface than allows implementing an INT 21h front-end hook. We'll look at some problems with the INT 21h redirector interface a bit later on.

Another way to attach foreign file systems to DOS is, of course, to use an installable block device driver. But the device driver interface imposes certain restrictions on the file system. It must have a `FILE DPE` and `FILE` which is not appropriate to a file system. There is the added fact that device drivers are a complex responsibility, though `DEVLOD` from Chapter 7 helps a lot!

The alternative to hooking INT 21h or writing a block device driver is to use the INT 21h function 11h redirector interface. It is often stated that "the network redirector grabs file system calls for the remote file system (INT 21h sees them) or words to that effect." While that is a good description of how the workstation-side `NTFV` works in `NetWare 2` and `3`, it is not a good description of the redirector, which operates at a level *below* or *inside* INT 21h. As we'll see, `MS DOS`'s code for INT 21h takes care of calling INT 21h Function 11h when appropriate.

**What DOS Provides** On one side of the redirector interface are the redirector services. These consist of DOS file structures and sets of INT 21h calls issued at strategic times by the DOS kernel. While the standard model of an operating system is that of a lower-level program that responds to requests for services initiated by an application, the redirector interface specifies function calls that the operating system provides and that a redirector must intercept and service. In normal DOS program operation, `INT 21h`, `ME`, and `IOFS` on the host program has hooked INT 21h services the call. In the case of network `DOS` calls, `INT 21h`, `ME`, and `IOFS` on the redirector service by `call Entry` (see `FILE` entry multiple) are explicitly set aside for redirector services, and each redirector function call will be a series of function calls, use `DOS` calls conventionally from inside INT 21h, such a program is a `block and redirector`, in contrast to the `NTFV` style front-end hook.

The data structures associated with the redirector interface include `SVASys`, the Current Directory Structure `CDS`, the System File table `SF1`, and the Swappable Data Area `SDA`.

The `SVASys` structure, obtained through undocumented INT 21h function 52h, provides the address of the `CDS` table and the `LASTDRIVE` value. A redirector responsible for initializing, maintaining, and if it is desirable, restoring the `CDS` for its own share of entries. For the most part, only deals with requests to the `CDS`, such as redirected `INT 21h` `ME` 11h. In contrast from the `IOFS` or `grd` subroutines before the reach the `CDS` system `Novell NetWare 2` and `3`, for example, can have `block and redirector` `LASTDRIVE`. `DOS` primarily cares about three fields in the `CDS`: the current directory string, the offset of the root directory entry string, and the flags word.

As you saw in the earlier sections of this chapter, a `SF1` entry holds the state that `DOS` maintains for each open file in a system. `DOS` appears to be interested only in a subset of the `SF1`: the open mode flags, which describe the access level to and shareability of the file; the device information flags word, which indicates whether the device is a block or character level device, whether it has been written to, a write failure, the date, time, and file size, and the current position fields. While the `SF1` is a `DOS` structure, its data fields are not for the use by the drive's owner. And while the `SF1` was designed for `DOS` to manage files in its `FAT`-based file systems, and so contains fields that deal with sets of sector and cluster, a redirector based on some other file system may use those fields in whatever way it needs.

The most important structure for the redirector interface is the Swappable Data Area. DOS is usually run in real mode, but it can be re-entered using the SDA. This is explained in more detail in Chapter 9. This section includes a section on building TSRs with the SDA. The SDA is that part of the DOS data segment which must be saved and restored to provide for DOS re-entrancy. More than that, however, the SDA is also that part of the DOS data segment that contains the critical data required for components to run a redirector. This data includes key IRPs (handles) such as the current critical error state, a pointer to the current Disk Transfer Area, the current open disk, to a file, search attributes for Find First, Find Next, and so on.

## Origins of the SDA

*Tim Farley*

**[Some programmers wonder why a single-tasking operating system like MS-DOS has a Swappable Data Area in the first place. The following brief historical note was written by Tim Farley, author of the forthcoming book *Undocumented NetWare*, to be published by Addison-Wesley in early 1994.]**

I always find it helpful to understand the historical background behind APIs. The fact that the "server" portion of NetWare Lite uses the DOS SDA got me to thinking about why DOS includes an SDA in the first place.

That NetWare Lite is a redirector provides a clue. Perhaps the SDA was part of the other Microsoft Networks code introduced with DOS 3.1 and was originally designed to allow the implementation of a background network server program, which would obviously get into some situations where it would need to re-enter DOS to serve multiple tasks. MS-NET and its derivatives like the IBM PC LAN Program were peer-to-peer networks very similar in concept to NetWare Lite.

To test this theory—that Microsoft introduced the SDA into DOS to support redirectors, and especially peer-to-peer "servers"—I went to Atlanta's Computer remnants and job lot outlet store called (oddly enough, "Quality Computer Components") and lo and behold, there was a copy of the IBM PC LAN Program 3.00, which I promptly bought for a grand total of \$2.50. (I believe this was one of the first, if not the first, implementation of Microsoft Networks (MS-NET) and therefore might even be the first redirector.)

I couldn't get the product (dated March 1985) to run on my hardware, but I was able to shoot it through the Source disassembler from V Communications, and sure enough, the server portion of IBM PC LAN 3.00 uses the DOS SDA function call (INT 21h AX=5D06h). This is RECEIVER.COM, similar to SERVER.EXE in NetWare Lite. After calling INT 21h AH=5D06h to get a pointer to the SDA, the code does a REP MOVSB, so you can see why the function uses the CX register to return the size of the swap IN DOS area.

As one would expect, the client portion (REDIR.EXE, the equivalent of CLIENT.EXE in NetWare Lite) implements a redirector and makes very extensive calls to twenty-seven different DOS internal (INT 2Fh AH=12h) functions. Clearly, the INT 2Fh AH=12h calls to the SDA were intended for redirectors.

REDIR.EXE, RECEIVER.COM, and NET.COM also all extensively use INT 2Ah, including critical sections 5 and 0fh.

IBM PC LAN Program 3.00 dates from March 1985, so that's the "date to beat" to find an earlier user of the DOS SDA.

When, as in disktasks, generally just copy entire blocks of the SDA without caring as to their contents (see Chapter 9), network redirectors must use specific fields in the SDA. Unfortunately, the SDA structure changed from DOS 3.0 to DOS 4.0, so programs that want to work in DOS 3.0 as well as in DOS 4.0 must make do with two different SDA structures. The Appendix shows the DOS 3.0 style SDA under INI 21h AX 5100h and the DOS 4.0 style SDA, which also applies to DOS 5.0 and 6.0 under INI 21h AX 5100h. Note, however, that the AX 5100h call is for DOS 4.0 *only*. Thus, in DOS 3.0 and 6.0, you use the same INI 21h AX 5100h call as in DOS 4.0, but if the SDA has the same format as a DOS 4.0. Totally confused yet? Fortunately, in DOS 4.0, you can also use INI 21h AX 5100h, so it's best for those writing redirectors to just forget about the AX=5100h call. Besides checking the DOS version to determine whether a DOS 3.0 or a DOS 4.0 style SDA is in use, a program can also consult a WORD flag in a DOS keep at offset 4 bytes data segment. For example, this shows MSCDEX determines which style DOS segment is in use:

```
CURR_DTA_PTR    dw 20Ah      ; offset in DOS 3.0S of CURR_DTA
CURR_DRV_PTR    dw 2E6h      ; offset in DOS 3.0S of CURR_DRV
; etc. other DOS variables

mov ah, 52h
int 21h
mov word ptr DOS_DS, es      ; get SysVars ptr, ignore BX; ES = DOS DS
mov ax, word ptr es [4]      ; DOS_DS[4] = style; 0 = DOS 3.0; 1 = DOS 4+
mov DOS_DS_VERS, ax

;
cmp DOS_DS_VERS, 1          ; is this DOS 4+ style DOS DS?
jz DOS_4_DS
jmp short done              ; What about values other than 0 or 1?

;
DOS_4_DS
mov CURR_DTA_PTR, 32Ch      ; offset in DOS 4+ DS of CURR_DTA
mov CURR_DRV_PTR, 356h      ; offset in DOS 4+ DS of CURR_DRV
done
```

Of course, that MSCDEX is *not* using the SDA, but instead contains two sets of offsets into the DOS data segment. MSCDEX uses DOS\_DS[4] to determine which set of offsets to use. Probably only Microsoft could afford to store inventories of use and word offsets into the DOS data segment. In general, third-party redirectors should use offsets from the start of the SDA address that INI 21h AX 5100h returns.

Of the many fields in the SDA, a redirector needs access to only a large handful:

ERR_LOCUS_ERR_CODE	Locus and extended error code of last error
ERR_ACTION_ERR_CLASS	Size, sense, action and class of last action
CURR_DTA	Current Disk Transfer Area (DTA)
FN1, FN2	Filename work areas
SRC_H, D_K	1 and 1st, 1st and Next search data block (SDB)
FOUND_H, F	Directory entries for found file
FN1_H, H1AR	File style filename file device name comparison
FN2_H, H2AR	File style filename file device name and destination
SRC_H, ATTR	Directory search attributes
OPEN_ATTR	File open mode
DRIVE_CDSMPTR	Pointer to CDS for drive accessed
SPOP_ACT	Extended open 21/0C action code
SPOP_ATTR	Extended open attributes
SPOP_MODE	Extended open mode
REN_SRC_H, I	Rename source search data block (SDB)
REN_H, I	Rename source directory entry



The DOS kernel lists a many of these fields as it services INT 21h calls. In fact, these are 90 fields in a generic structure, but variables in DOS's own data segment, to which redirectors are given (or at least take) privileged access via the SDA.

Listing 8-24, DOSSTRUC.SCR, is an INTRSPY script that describes only the fields in the DOS structures that are important to a redirector. Other script files used to investigate the interface will include this script. The script uses a feature of INTRSPY that we failed to mention in Chapter 5 (Undo, correct INTRSPY). The ability to specify fields in a structure that INTRSPY skips over when asked to dump the entire structure.

#### Listing 8-24: DOSSTRUC.SCR

```
; DOSSTRUC.SCR
; Undocumented DOS structures relevant to the redirector interface
; This uses the undocumented INTRSPY "skip" feature.

;; Current Directory Structure entry - All DOS versions
structure CDS fields
  CURR_PATH (byte,asciiz,87)
  CDS_FLAGS (word,hex)
  FT (byte,skip,10)
  ROOT_OFS (word,dec)
  ; In DOS 4.0 and above there are a further 7 bytes of IFS/SHARE fields

;; Directory entry for 'found' file - All DOS versions
structure DIRENTRY fields
  FNAME_11 (char (byte,char,11))
  ATTR_BYTE (byte,hex)
  FD (byte,skip,10)
  F_TIME (word,hex)
  F_DATE (word,hex)
  START_CLUSTER (word,hex) ; A redirector can reuse this field
  F_SIZE (dword,dec)

;; Search Data Block - All DOS versions
structure SDB fields
  DRIVE_NUM (byte,dec)
  SRCH_MASK (byte,char,11)
  DIR_ENTRY (word,dec) ; A redirector can reuse this field
  DIR_SECTOR (word,dec) ; A redirector can reuse this field
  FD (byte,skip,4)

;; Lock/Unlock region of file structure
structure LOCKREC fields
  REGION_START (dword,dec)
  REGION_LEN (dword,dec)
  FD (byte,skip,13)
  FILENAME (byte,asciiz,80)

;; Swappable DOS Area - DOS 3.1 to 3.5
;; Used in the form SDA_SEG:SDA_OFS->SDA3
structure SDA3 fields
  FD (byte,skip,3)
  ERR_LOCUS (byte,hex)
  ERR_CODE (word,hex)
  ERR_ACTION (byte,hex)
  ERR_CLASS (byte,hex)
  DEVDRVR_PTR (dword,ptr)
  CURR_DTA (dword,ptr)
  FT (byte,skip,30)
  DD (byte,dec)
  MH (byte,dec)
  YY_1980 (word,dec)
  IZ (byte,skip,96)
  FN1 (byte,asciiz,128)
  FN2 (byte,asciiz,128)
```

```
SRCH_BLK (SDB)
FOUND_FILE (DIRENTRY)
F3 (byte,skip,81)
FN1_11CHAR (byte,char,11)
F4 (byte,skip)
FN2_11CHAR (byte,char,11)
F5 (byte,skip,11)
SRCH_ATTR (byte,hex)
OPEN_MODE (byte,hex)
F6 (byte,skip,48)
DRIVE_CDSPTR (dword,ptr)
F7 (byte,skip,72)
REN_SRCFILE (SDB)
REN_FILE (DIRENTRY)
```

;; Swappable DOS Area - DOS 4.0, 5.0, 6.0

;; Used in the form SDA\_SEG:SDA\_DFS->SDA4

```
structure SDA4 fields
FD (byte,skip,3)
ERR_LOCUS (byte,hex)
ERR_CODE (word,hex)
ERR_ACTION (byte,hex)
ERR_CLASS (byte,hex)
DEVDRVR_PTR (dword,ptr)
CURR_DTA (dword,ptr)
F1 (byte,skip,32)
DB (byte,dec)
MH (byte,dec)
Y1_1980 (word,dec)
F2 (byte,skip,106)
FN1 (byte,ascii,128)
FN2 (byte,ascii,128)
SRCH_BLK (SDB)
FOUND_FILE (DIRENTRY)
F3 (byte,skip,88)
FN1_11CHAR (byte,char,11)
F4 (byte,skip)
FN2_11CHAR (byte,char,11)
F5 (byte,skip,11)
SRCH_ATTR (byte,hex)
OPEN_ATTR (byte,hex)
F6 (byte,skip,51)
DRIVE_CDSPTR (dword,ptr)
F7 (byte,skip,87)
SPOP_ACT (word,hex)
SPOP_ATTR (word,hex)
SPOP_MODE (word,hex)
F8 (byte,skip,29)
REN_SRCFILE (SDB)
REN_FILE (DIRENTRY)
```

;; System File Table entry - All DOS versions

```
structure SFT fields
C_HANDLES (word,dec)
OPEN_MODE (word,hex)
ATTR_BYTE (byte,hex)
DEV_INFO (word,hex)
DEVDRVR_PTR (dword,ptr)
S1_C_STR (word,dec) ; A redirector can reuse this field
F_TIME (word,hex)
F_DATE (word,hex)
F_SIZE (dword,dec)
F_POS (dword,dec)
LAST_RELC_STR (word,dec) ; A redirector can reuse this field
LAST_ABSL_STR (word,dec) ; A redirector can reuse this field
DIR_SCTR_NO (word,dec) ; A redirector can reuse this field
```

```
DIR_ENTRY_NO (byte,dec) , A redirector can reuse this field
FNARE_11CHAR (byte,char,11)
; other fields
```

Note the “A redirector can reuse this field” comment at some places in `DDOS\RUUCSR`. When we discuss the Phantom source code, you will see that Phantom uses a slightly different version of the NEI than that shown earlier in this chapter. The meaning of the fields that DDOS may reuse must of course be left unchanged. Similarly, DDOS expects that the redirector will fill in certain fields. This expectation is a critical part of the redirector interface.

**What a Redirector Must Supply** On the other side of the interface is the redirector itself. In Microsoft's (the PC LAN) program, the hypothetical CompuServe file system, or in this chapter's demonstration redirector, the Phantom. A redirector normally subsists as a TSR program and to install into the chain of INE 21b handlers, gets the vector to the current INE 21b handler, stores it as the next handler in the chain, and then sets the INE 21b vector to point to itself. When INE 21b is invoked, the redirector's INE 21b handler receives control. If the INE 21b call is for the redirector, `API_11b` will pass control to the next handler in the chain. In this way, the redirector monitors all INE 21b calls and intercepts all but redirector interface actions only.

One major problem with this interface is, since INE 21b always starts with the INECHAIN program, CompuServe's (and other) TSR and other may be competing for this multiplex interrupt. It is not uncommon to have ten or more handlers in the INE 21b chain (see chapter 11). The performance problems with INE 21b may be one possible reason for Microsoft's apparent move away from the redirector interface—the direction of a Novel NEA vs. a local call, look, even when Novel itself is moving toward the redirector.

Another problem with the redirector interface is that it provides no assistance when multiple redirectors are present. Each redirector must decide on a subfunction by subfunction basis which INE 21b API calls call it, which chain to another, previously installed redirector. We discuss MS-DOS and NetWare's use of the redirector interface, the other interface, in section 8.6.

Because INE 21b redirectors are called by DDOS, they have no key bytes, and a way for INE 21b redirectors which grab I/O requests before DDOS even sees them (for example, INE 21b redirectors have little access over the DDOS file naming scheme, only `server\0025`). However, whereas an INE 21b hook at least theoretically can delay any file naming, copy, or other activity.

### Tracing an Open, Revisited

Much earlier in this chapter, we used `INTRSPY` to trace through the `COPY` command working on a local drive. `COMMAND.COM` called INE 21b API 16b, which in turn called a block device driver which then (as we see in figure 8-4) called INE 13b. We now must support an exercise for a redirector drive. This can be a network drive, a CD-ROM (as provided by MS-DOS), or even an MS-DOS disk provided by the Phantom sample program described later in this chapter.

We can use the exact same `DISKSRV` (figure 8-4) used to produce figure 8-1, because `DISKSRV` includes DDOSR and network redirector drivers don't have associated block device drivers. However, removing DDOSR results in an `INTRSPY` script that we can use to see what happens to a `COPY` command on a redirector drive.

```
C:\NDDOC2>phantom -s256 h:
256kb XMS allocated
Phantom installed as H
C:\NDDOC2>copy con h:foo bar
this is foo.bar
^Z
      1 file(s) copied
C:\NDDOC2>intrspy
C:\NDDOC2>cdspy compile disk command /c copy h:foo bar h:bar.foo
```

However, the resulting INTRSPY report doesn't show anything occurring "inside" the INT 21h function 6Ch call:

```
21/6C: Ext Open/Create: H:FOO.BAR
AX=6C00 BX=0040 CX=0000 DX=0101
21/6C: done, file is 0005
```

We i of course it doesn't show anything inside the INT 21h AH=6Ch call! Why should it? There is no block device driver installed for the drive, and it seems unlikely that an XMS RAM disk would call INT 13h, which is the BIOS disk function.

So what *do* happen inside an INT 21h on a network redirector drive? DOS issues INT 21h AH=11h calls to the network redirector. To see this, we can create a very simple INTRSPY script which tracks entry and exit into a few INT 21h calls and tracks any calls to INT 21h AH=11h. 2F11.SCR is shown in Listing 8-25.

### Listing 8-25: 2F11.SCR

```
2F11.SCR
Intercept 21h
function 3Ch
on_entry output "21/3C: Create (ds:dx->byte,es:ix,64)
on_exit output "21/3C: Create done: " ax
function 3Dh
on_entry output "21/3D: Open " (ds:dx->byte,es:ix,64)
on_exit output "21/3D: Open done: " ax
function 3Eh
on_entry output "21/3E: Close " bx
on_exit output "21/3E: Close done " bx
function 3Fh
on_entry output "21/3F: Read " bx
on_exit output "21/3F: Read done " bx
function 40h
on_entry output "21/40: Write File " bx
on_exit output "21/40: Write done " bx
function 6Ch
on_entry output "21/6C: Ext Open/Create " (ds:xi->byte,es:ix,64)
on_exit output "21/6C: Ext Open/Cr done: " ax
Intercept 21h
function 11h
on_entry output " 2f/" ax
run "command /c %1 %2 %3 %4 %5 %6 %7 %8 %9"
report
```

It's a little more interesting than the script in Listing 8-24.

```
C:\UNDOC2>cdspy compile 2f11 copy h:foo bar h:bar:foo
```

```
21/6C: Ext Open/Create H:FOO.BAR
2F/1123
2F/112E
21/6C: Ext Open/Cr done: 0005
21/3F: Read 0005
2F/1108
21/3F: Read done 0005
21/3E: Close 0005
2F/1106
21/3E: Close done 0005
21/6C: Ext Open/Create H:BAR.FOO
2F/1123
2F/112E
21/6C: Ext Open/Cr done: 0005
```

```

2F/1123
21/40: Write File 0005
2F/1109
21/40: Write done 0005
21/3E: Close 0005
2F/1106
21/3E: Close done 0005

```

The ordering of INT 2fh calls within INT 2fh calls provides a clear picture of how the network redirector interface works. For example, calling INT 2fh AH=3fh (Read File) on a network redirector drive results in DOS sending an INT 2fh AX=1108h to the redirector.

```

21/6C Open/Create -> 2F/1123 and 2F/112E
21/3F Read -> 2F/1108
21/40 Write -> 2F/1109
21/3E Close -> 2F/1106

```

In DOS 5.0 and higher, the COPY command happens to use INT 2fh AH=6fh. If a file had instead been opened with the older INT 2fh AH=3fh call, DOS would have issued an INT 2fh AX=1116h rather than the AX=112Eh seen above. It is important to understand that DOS usually only makes these 2F-11 calls when the network bit is set in the current CDS's flags or in an SFT entry's device\_info word. For example, in a flow MS-DOS decides whether to issue the INT 2fh AX=1108h Read call

```

; ES:DI points to an SFT entry
test byte ptr es [di+6], 80h ; test network bit in SFT dev info
jz not_redirect
mov ax, 1108h
int 2fh

```

not\_redirect:

Likewise, here's how DOS decides whether to make the INT 2fh AX=1105h Mkdir call

```

; ES:DI points to the current CDS
test byte ptr es [di+44h], 80h ; test network bit in CDS flags
jz not_redirect
mov ax, 1105h
int 2fh

```

not\_redirect:

Even if the network bit in the CDS says the network bit is set, there are still a few redirector calls that DOS issues. Experimentation with INTRSPY and inspection of a DOS disassembly show that several redirector calls are unique situations: 10h, 20h, 22h, 23h, and 25h. In essence, bit 0advsts that DOS issues, e.g. to check whether a redirector is installed or not.

Furthermore, there are a few relatively obscure redirector calls, sub-sections 18h and 19h, that don't require a CDS entry. These are used, for example, by FAN Manager named pipes. But for the most part, starting a redirector means setting the network bit in a CDS entry. Likewise, a redirector must set the network bit in the SFT device\_info word for any files opened on its drives.

So far, it's not like there are just a few redirector calls, and that these correspond closely to the execution of INT 2fh calls. However, we haven't done much to the redirector drive yet - just copy a single call to the Windows registry (DIF, DIR, MD, CD, RD, and so on) to get a general idea of which redirector calls DOS makes. We can simplify the INTRSPY script even further to print out nothing but the INT 2fh AX function number

```

intercept 2fh function 11h
on_entry output ax

```

If we then go a few minutes work on the drive and sort the results, we can get some idea of the range of redirector calls that DOS issues.

```
C:\UNDOS2>cmdspy report | sort | uniq
```

```
1101      ; remove dir
1103      ; make dir
1105      ; change dir
1106      ; close file
1108      ; read file
1109      ; write file
110C      ; get disk space
1113      ; delete file
1116      ; open file
1117      ; create/truncate file
1118      ; find first
111C      ; find next
111D      ; close all for PSP
1120      ; flush all buffers
1122      ; proc term hook
1123      ; qualify pathname
1125      ; redirected printer mode
112E      ; extended open file
```

And yes, you can detect a lot of random activity on the drive. Even more, there are other, less common but still important redirector calls that DOS makes. For example, calling INT 21h AX=5d02h sets the NETWORK process. Listing 8-15 results in a call to INT 21h AX=111fh. Basically, it sets up a procedure table redirector requires implementing the whole blocks interface. You can do a lot more. For example, you just implemented the calls listed above.

Capturing calls with INTRSPY is a fairly random way to determine what parts of the redirector interface are used. To see what redirector calls DOS makes, it is better to examine the actual DOS interface so as to support. Using the INT21.FST file generated with NCDDBG in Chapter 6, you can see exactly which INT 21h AH=11h calls DOS makes and when. For example,

```
INT21_0B:      ; Reset Drive

; do INT 2Ah AX=800th, call func to flush all buffers (calls code for
; INT 2fh AX=1215h in loop), do INT 2Ah AX=810th

FDCB:4DDC  0BFFFF      MOV AX,FFFF
FDCB:4DDF  50           PUSH AX
FDCB:4DE0  0B2011      MOV AX,1120      ; flush all disk buffers
FDCB:4DE3  CD2F       INT 2F           ; 21/1120
FDCB:4DE5  58           POP AX
FDCB:4DE6  C3           RET
```

Inspecting this DOS 5.0 kernel indicates that it calls redirector subfunctions 0, 1, 3, 5, 0Ah, 0Ch, 0E, 0Fh, 11, 13, 16h, 19h, 1B, 20h and 2Fh. Therefore, a commercial redirector should implement all these DOS calls, which are described in a moment.

So far we have seen which redirector calls DOS makes but haven't seen any details of how they work. To see the interface in more detail, we can follow a file Open call a little more closely using another INTRSPY script, REDIR.SCR. Listing 8-26.

### Listing 8-26: REDIR.SCR

```
;; REDIR.SCR == for DOS >= 4.00
Include "dosstruc"      ; see listing 8-24
Intercept 21h
function 0Ch
on_entry
    output      DOS_OPEN (6ch) "XXXXXXXX-XXXX-XXXXXX"
    output "File name: " (ds:si->byte,mac16,40)
    output "Open mode: " (bl,hex) "h"
```

```

on_exit
  output
  output "== 6Ch OPEN Completed "
  if (cflag == 1) sameLine "(FAILED " (ax,dec) ") ====="
  if (cflag == 0) sameLine "(Handle " (ax,dec) ") ====="

```

```
Intercept 2fh
```

```
function 1fh
```

```
subfunction 2Eh
```

```
on_entry
```

```

output
output "-- 2F Open (2Eh) -----"
output "File name " (sda_seg:sda_ofs->SDA4,FN1)
output "Open mode: " (sda_seg:sda_ofs->SDA4,SPOP_MODEE)
output "Uninitialized SFT:"
output (es:di->SFT)

```

```
on_exit
```

```

output
output " 2F Open (2Eh) completed "
if (cflag == 0)
  sameLine "-----"
  output "Completed SFT:"
  output (es:di->SFT)
if (cflag == 1) sameLine "(FAILED " (ax,dec) ") -----"

```

```
run "command /c type %1"
```

```
report
```

This script takes as a parameter the name of a file that exists on the redirected drive. For variety, we'll look at a file on an MSCDEX (CD-ROM) drive, rather than on Phantom drive H.

```
C:\UNDOC2>cmdspy compile redir l:read me > redir log
```

This types out the file entered as the parameter. If the file is on a redirected drive, as in the MSCDEX-based example, `READ ME`, the `REDIR LOG` contains something like this:

```

== DOS OPEN (6Ch) =====
File name: l:read me
Open mode: 0Dh

-- 2F Open (2Eh) -----
File name: \\.\A:\READ.ME
Open mode: 0000
Uninitialized SFT

SFT_C_HANDLES           : 65535
SFT_OPEN_MODE           : 0000
SFT_ATTR_BYTE           : 20
SFT_DEV_INFO             : 0042
SFT_DEVORV_PTR           : 0116:13A4
SFT_ST_CLSTR             : 7002
SFT_F_TIME               : 2800
SFT_F_DATE               : 1689
SFT_F_SIZE               : 47845
SFT_F_POS                : 47845
SFT_LAST_RELCLSTR       : 5
SFT_LAST_ABSCLSTR       : 867
SFT_DIR_CLSTR_NO        : 0
SFT_DIR_ENTRY_NO        : 11
SFT_FNAME_T1CHAR         : COMMAND.COM

-- 2F Open (2Eh) completed -----
Completed SFT

SFT_C_HANDLES           : 65535
SFT_OPEN_MODE           : 0002
SFT_ATTR_BYTE           : 01
SFT_DEV_INFO             : 8048

```

```

SFT_DEVDRV_PTR          : 0057 00F3
SFT_ST_CLSTR            : 9933
SFT_F_TIME              : 2395
SFT_F_DATE              : 7ADC
SFT_F_SIZE               : 21000
SFT_F_POS                : 0
SFT_LAST_RE_CLSTR      : 21560
SFT_LAST_ABSCLSTR      : 0
SFT_DIR_SCTR_NO         : 56026
SFT_DIR_ENTRY_NO       : 2
SFT_FILENAME           : READ ME
** 6Ch OPEN Completed (Handle 5) =====

```

This shows the DOS Open function called with the raw filename string in DS:SI as specified in reading 2 and the resultant redirector Open call subfunction 16h called with the SFA:SI field directly a fully qualified filename: \\FAR\MEM. Note that the filename for files on network redirected filesystems begins with a double backslash. It also shows an unformatted SFI passed to the redirector Open function, with data left over from previous use. COMMAND.COM in this case used the SFI initialized by MSCDEX. This gives you an idea of the sort of detail DOS takes care of before calling a redirector, as well as the tasks for which the redirector is responsible.

### The Phantom

To look for more detail, it is best to use a real example. Here, then, is the Phantom. The Phantom is perhaps a phantom drive, that is, the drive is not based on a physical disk drive, but rather uses VMS as its media. Radically changed from its guise in the first edition of *Undocumented DOS* as the world's least effective storage device, Phantom is now a fully functioning and useful extended memory RAMDisk TSR.

- It supports all DOS file system commands that you can run on a network drive, with the exception of CHKDSK, FORMAT, SYS, and DELTREE. These all operate on native DOS devices such as 5.25- or 3.5-inch 5.25- and 7.25-inch floppy disks, which do not get redirected through the redirector interface.
- It is even faster than a very fast hard disk and is transparent to all DOS and Windows applications. Note, however, that loading SHARE has no effect. Phantom does not signal sharing violations.
- It works under DOS 3.00 to 7.00, exclusive. It has also been tested under a prerelease version of DOS 7.00 (beta).
- It supports CBI as well as handle-based file operations.
- Unlike a first version of Phantom, the number of directories and files on the drive is limited only by the amount of VMS memory allocated for the drive on the command line. As this book was going to press we found that Phantom recently supports only two subdirectory levels.
- As an extended memory RAMDisk redirector, it has some advantages over conventional device driver RAM disks like RAMDRIVE.SYS (1) you can load it when you need it, unload it when you don't need it, and reload it onto a VMS allocation on the fly to suit the application.
- It illustrates not only a shared and unshared knowledge of the redirector interface, but principles of a FAT-based file system as well.

Its usage is

```

PHANTOM [-Snnnn] d: [-U]
-Snnnn  specifies size of RAM disk in KB of EMS
d:      specifies drive letter to use
-U      unloads latest copy of Phantom loaded

```



The `SWITCH` is particularly important because by default the Phantom takes all XMS memory, this is not necessarily what you want.

You can run multiple copies of Phantom, specifying a different drive letter each time. Note that `C` unloads the most recent invocation.

Here's a brief sample session with the Phantom drive, where the entire contents of the Interrupt List from `C:\NDOS\INTRLIST` is XCOPIED to a new `INTRLIST` subdirectory on Phantom drive `H`:

```
C:\NDOS2\CHAP8>phantom h:
396Kb XMS allocated
Phantom installed as H:

C:\NDOS2\CHAP8>md h:\intrlist
C:\NDOS2\CHAP8>xcopy C:\NDOS2\INTRLIST\*. * h:\intrlist
Reading source file(s)...
C:\NDOS2\INTRLIST\COMBINE.BAT
C:\NDOS2\INTRLIST\FILE_ID.DIZ
C:\NDOS2\INTRLIST\GLOSSARY.LST
C:\NDOS2\INTRLIST\INTERRUP.1ST
C:\NDOS2\INTRLIST\INTERRUP.A
Reading source file(s)
C:\NDOS2\INTRLIST\INTERRUP.B
; ... okay, we got the idea.
C:\NDOS2\INTRLIST\*_ADVERT.TXT
30 file(s) copied

C:\NDOS2\CHAP8>dir h:\intrlist
Volume in drive H is PHANTOM
Directory of H:\INTRLIST

COMBINE.BAT          150 05-06-93   11:06a
FILE_ID.DIZ         200 05-24-93   12:34a
GLOSSARY.LST       23998 03-21-93    4:42p
INTERRUP.1ST       60877 05-24-93   12:34a
INTERRUP.A         290154 05-24-93   12:34a
INTERRUP.B         242488 05-24-93   12:34a
; ... lots of big files
*_ADVERT.TXT       1845 01-28-92   10:19p
30 file(s)          3096942 bytes
                   941056 bytes free
```

Naturally, processing the three-megabyte Interrupt List on this Phantom drive, whether with a tool like `dir` or with Hackswack Software's `INTRVAL`, is very fast.

All file entries (except if known file systems are here) have created an entry which looks like a drive. Behaves like a drive, stuff which lives outside our IN12H In-Station File handler and XMS. As this book was going to press, a beta site discovered that Phantom crashes if you have `HMMMSYS` (or don't have `DDOS=HIGG`). This right now Phantom inadvertently requires `DDOS=HIGG`.

### Phantom Implementation

With the C source code for Phantom, we can resume our trace into the `DDOS` Open, call and see how a redirector like Phantom handles this and other `IN12H` calls that `DDOS` sends it. The source code for Phantom is over 1,740,000 lines of code, too print out in its entirety, but if you're accompanying disk as `C:\NDOS2\CHAP8\PHANTOM.C`. Here, we look at selected portions of `PHANTOM.C` to see how a redirector handles `DDOS` Open, Read, Find First, Change Directory, and Make Directory.

**Initializing the CDS** First, however, we must see how the Phantom sets itself up to be called in the first place. As we saw earlier, `DDOS` issues most of the `IN12H` API File handler redirector calls only if

the drive letter is set in the current CDS or in the device info word of the current SEI. Before you can see how the Phantom handles INT 13h calls, then you must see how it sets up the CDS. This, appropriately enough, is done in PHANTOM.C's set\_up\_cds() function shown in Listing 8.27.

### Listing 8.27 PHANTOM.C set\_up\_cds()

```
// V3_CDS_PTR and V4_CDS_PTR are CDS structures for DOS 3, DOS 4+
// LOLREC_PTR is pointer to LOLREC (LOL = List of Lists = SysVars)
// global variables
LOLREC_PTR lolptr, /* pointer to List Of Lists */
uchar our_drive_no, /* A is 0, **not* 1 !! */
uint cds_root_size, /* Size of our CDS root string */
char far * cds_path_root = "Phantom \\", /* Root string for CDS */
extern void failprot(char * msg); /* print message, exit to DOS */
void set_up_cds(void)
{
    V3_CDS_PTR our_cds_ptr = lolptr->cds_ptr;
    if (our_drive_no >= lolptr->last_drive)
        fa (prot:"Drive letter higher than last drive ");
    if (_osmajor == 3)
        our_cds_ptr += our_drive_no;
    else
    {
        V4_CDS_PTR t = (V4_CDS_PTR) our_cds_ptr,
        t += our_drive_no,
        our_cds_ptr = (V3_CDS_PTR) t;
    }
    // Check that this drive letter is currently invalid (not in use)
    // Yes! both Physical (0x4000) and Network (0x8000) at same time.
    if ((our_cds_ptr->flags & 0xc000) != 0)
        fa (prot:"Drive already assigned...");
    // Establish our 'root'
    cds_root_size = _fstrlen(cds_path_root);
    _strcpy(our_cds_ptr->current_path, cds_path_root);
    our_cds_ptr->current_path[_fstrlen(our_cds_ptr->current_path) - 3] =
        (char) ('A' + our_drive_no);
    _strcpy(our_cds_ptr->root, our_cds_ptr->current_path);
    our_cds_ptr->root_off = _fstrlen(our_cds_ptr->current_path) - 1;
    current_path = our_cds_ptr->current_path + our_cds_ptr->root_off;
    // Set both Physical (0x4000) and Network (0x8000) at same time.
    // Dave D Markov says that a non-network redir such as Phantom
    // should set 0xc080, not 0xc000. 0x80 is the REDIR_NOT_NET bit
    // used by MSCDEX. If 0x80 is not set, Phantom shows up in
    // NET USE in DEC Pathworks and IBM LAN Server. This point
    // comes from an Interrupt List entry supplied by Geoff Chappell.
    our_cds_ptr->flags |= 0xc080;
}
```

On entry to set\_up\_cds(), the global variable our\_drive\_no holds the zero-based A-FH drive letter specified for the new Phantom drive, and set\_up\_cds() uses this to get a far pointer to the CDS entry for that drive. The function ensures that the drive doesn't exceed F:\A:\DRIVE (in other words, that the CDS is large enough) and ensures that the drive is not already in use by testing both for Physical and Network bits in our\_cds\_ptr->flags.

Assuming that the specified drive letter can be Phantomized, the function sets the current\_path in the CDS entry to a default string such as "Phantom H" and sets the root\_off field to the backslash, so that the root directory is ". For this drive, the TRUENAME command would print the string "Phantom H" directly from the CDS.

Most important, set up `cds` turns on the CDS gate's Network and Physical bits (our `cds_ptr_flags = 0x0000`) thereby declaring that disk, directory, and file requests for this drive will be handled by a redirector. If you remove a Phantom drive with PHANTOM.U, the program switches off these bits (our `cds_ptr_flags &= 0x0000`). See the comment in `set_up_cds`, which notes that this should really be `0x0080` not `0x0000`. As you can tell, this was a last-minute change.

**The Redirector INT 2Fh Handler** Once the appropriate bits have been set, there must of course, be an INT 2Fh handler to act as the redirector. The Phantom installs its INT 2Fh handler, called `redirector()`, in the normal way:

```
prev_int2f_vector = _dos_getvect(0x2f); // 2f/35
_dos_setvect(0x2f, redirector); // 2f/25
_dos_keep(0, tar_params); // 2f/31: TSR
```

The Phantom's INT 2Fh handler is shown in Listing 8-28.

### Listing 8-28. PHANTOM.C `redirector()`

```
// All_REGS are computer-specific register params to interrupt functions
#define MAX_Fxn_NO 0x2E
#define STACK_SIZE 1024
#define FCARRY 0x0001

// global variables
int curr_fxn, /* Record of function in progress */
ALL_REGS, /* Global save area for all caller's regs */
uint dos_ss, /* DOS's saved SS at entry */
uint dos_sp, /* DOS's saved SP at entry */
char our_stack[STACK_SIZE], /* our internal stack */
uint far* stack_param_ptr, /* ptr to word at top of stack on entry */
int filename_is_char_device, /* generate_fcbname found CHAR dev? */

void succeed(void) { r_flags &= FCARRY, r_ax = 0, }

// dispatch table, fcnmap, and is_call_for_us() are discussed below
void interrupt far redirector(ALL_REGS entry_regs)
{
    static uint save_bp,

    __asm $f1,

    // Make sure that AH ish and that AL <= MAX_Fxn_NO
    // i.e., only support INT 2Fh AH=1100h through AH=1125h
    if ((entry_regs.ax >> 8) != (uchar) 0x11) |
        ((uchar) entry_regs.ax > MAX_Fxn_NO)
        goto chain_on;

    curr_fxn = fcnmap[(uchar) entry_regs.ax], // A's subfunction
    if ((curr_fxn == unsupported) ||
        (! is_call_for_us(entry_regs.es, entry_regs.d))) // listing 8-29
        goto chain_on,

    // Set up our copy of the registers
    r = entry_regs,

    // Save ss sp and switch to our internal stack. We also save bp
    // so that we can get at any parameter at the top of the stack
    // (such as the file attribute passed to subfxn 17h)
    __asm mov dos_ss, ss;
    __asm mov save_bp, bp,
    stack_param_ptr = (uint far*) &_FP(dos_ss, save_bp + sizeof(ALL_REGS));
    __asm (
        mov dos_sp, sp,
        mov ax, ds
```

```

c:1
mov ss, ax          // New stack segment is in data segment.
mov sp, offset our_stack + STACK_SIZE - 2
sti
}

succeed(); // Expect success!
// DO IT! Call the appropriate handling function
// unless we already know we need to fail (NULL, (DN, AUR, etc.)
// fs:aname_is_char_device is set inside is_call_for_us()
if (!fs:aname_is_char_device) fail(5),
    else dispatch_table[curr_fxn](),
// Switch the stack back
_asm {
    cli,
    mov ss, dos_ss,
    mov sp, dos_sp;
    sti;
}
// put the possibly changed registers back on the stack, and return
entry_regs = r;
return,
// If the call wasn't for us, we chain on.
chain_on:
    chain_intr(prev_int2f_vector),
}

```

The `is_call_for_us` function sees all INT 21h calls, whether they are redirector-related or not, and which of these are intended for one of the Phantom drives or not. For every call

- if it is a redirector call, that is AH=11h, and
- if it is a call for one of its drive letters (see below), and
- if it is a supported call,

then `redirector` dispatches the appropriate procedure to carry out the requested subfunction. If any of these parameters is false, the redirector passes control to the previous handler in the (possibly very long) INT 21h chain.

Phantom also cautions that you do not try to create files on the Phantom drive with names such as `NUL:*.exe` (DOS, that being an unmodified character device drivers). This is call for `is_call_for_us` function calls `generate_fcb` to see whether `fname` is a character device, which is a C wrapper around INT 21h `AN 1225` `generate_fcb` entry, then sets the global variable `fname_is_char_device`. Incidentally, `generate_fcb` may also have other references to FCBs in the Phantom code, but nothing to do with FCBs. This very confusing name merely refers to 11 character file names, such as one finds in FCBs but also in many other parts of DOS.

Having seen it just that it has a supported redirector call for one of its drives, `redirector` prepares a copy of the registers as they were on entry and then switches SS:SP to an internal stack, which is simply a static array of bytes in the `SRB_DS` upon entry to `redirector`. SS is the DOS stack segment, but DS becomes our DS context. (The C interrupt keyword "switching to our own stack satisfies the C requirement.") Furthermore, the DOS stacks are not very large—about 300 bytes, see Chapter 6—and having our own stack ensures that we do not run out of stack space.

Note, generally, is the only concern that use of this internal stack might raise. By using an internal stack, a debugger positioned to the register on entry, our redirector is not recntrant—the contents of the stack are always destroyed by each use at on at on. However, our redirector is called only by DOS, and, as we saw in Chapter 6, DOS does exactly the same thing upon entry to `int 21h` for all the functions which generate calls to `is`. To all intents and purposes, DOS itself is not recntrant, so our lack of

recursively instead of `swi`. However, the reader may wish to ponder how this interacts with software that uses the Swapable Data Area to recenter DCS.

Historically, stacked stacks redirector dispatches the appropriate handler function for the requested subfunction. When the subfunction returns, redirector switches SSIP to their original values upon entry, restores the registers, and returns.

**How Do We Know the Call Is for Us?** DCS allows the installation of multiple redirectors. If `redirector = ...` receives an INT 2E or AH=11h call, but one which is not intended for one of its drives, the function passes the call on in the normal way with the `chain_ptr` function, provided with both Microsoft C and recent versions of Borland C++.

How does the redirector check that it should handle a call rather than pass it on to MSDOS? NetWare File System (the redirector) Stack Drive may be a chain of redirectors, each wanting to serve with those calls that relate to the drive(s) that it is redirecting; there must be a way of determining whether a particular redirector call that's making the rounds is for you.

The redirector function passes the FS and DD registers to `is_call_for_us`, for us to figure out, well, whether the call is for us. This function is shown in Listing 8-29.

### Listing 8-29 PHANTOM.C is call for us()

```
// _clsfil, _unlockfil, _inquiry, etc are 2F/11 AL subfunction numbers
// SFIREC_PTR for ptr to an System File Table (SFT) entry
// SRCHREC_PTR for ptr to a findFirst Search Data Block (SDB)
// V3_SDA_PTR, V4_SDA_PTR for ptrs to the Swappable Data Area (SDA)
int is_call_for_us(int es, uint ds)
{
    filename_is_char_device = 0;
    // The first 'if' checks for the bottom 6 bits of the
    // device information word in the SFT Values > LASTDRIVE
    // we are to file not associated with LANs, such as LAN Manager
    // named pipes (thanks to David Markun)
    if ((curr_fxn == _clsfil || curr_fxn == _unlockfil)
        || (curr_fxn == _skfend)
            || (curr_fxn == _unknown_fxn_2D)) // file related
    {
        SFIREC_PTR sft_ptr = (SFIREC_PTR) PK_FP(es, ds); // check SFT
        // Markun says 0x3f mask is WRONG! See subfunction 1Ch below
        return (sft_ptr->dev_info_world & 0x3f) == our_drive_no;
    }
    else if (curr_fxn == _inquiry) // 2F/11/00 succeed automatically
        return TRUE;
    else if (curr_fxn == _fnext) // find next
    {
        SRCHREC_PTR psrchrec; // check search record in SDA
        if (_osmajor==3)
            psrchrec=&(V3_SDA_PTR) sda_ptr->psrchrec;
        else
            psrchrec=&((V4_SDA_PTR) sda_ptr->psrchrec);
        // Markun says 0x3f mask is WRONG! See subfunction 1Ch below
        return (psrchrec->drive_no & 0x3f) == our_drive_no;
    }
    else // everything else
    {
        uchar far * p;
        if (_osmajor==3)
            p = (V3_SDA_PTR) sda_ptr->cdsptr; // check CDS
        else
            p = ((V4_SDA_PTR) sda_ptr->cdsptr);
        if (_memcmp(cds_path_root, p, cds_root_size) == 0)
            return TRUE;
    }
}
```

```

    // If a path is present, does it refer to a character device?
    if (curr_fxn != _diskio)
        generate (cbname(ds), // eventually call 2F/1223
        return TRUE,
        )
    else
        return FALSE,
    }
}

```

The return value for `is_dev` depends on the type of operation to be performed. If the 2F/11 call is successful (i.e., an open file such as Read, Write, Com out, or Close), DOS sets up FN DI to point to the NCR file table by issuing the FN DI block. See the redirector specification later in this chapter. The kernel information word in the NCR entry contains the file's drive number in the 0x00000000-0000000F range, and comparison operation for `is_dev` call for `is_dev` to ascertain that the call was one of ours:

```

*(DWORD_PTR) PR_FP(es,di) > dev_info_word & 0x3F) // our_drive_no

```

If the check for subtraction is "Field Next" (the SRCH link structure in the NDA contains a drive number that is not a byte in the structure) and again we need only mask and compare the bottom 6 bits of that drive to see if the index is for us:

```

*(char *) es:cbec > drive_n & 0x3F) // our_drive_no

```

Actually, this 0x3F mask should probably be changed to 0x1E. In addition, bit 6 (0x40) should be `*(char *) es:cbec > 0x40`. Later in this chapter, for the gory details. Another last minute correction by David Markus!

In the case of the DOS window, already pointed to `DRIVE_CDS` (DWORD) field of the NDA at `0x00000000`, the drive is bit 2 (access during the redirector call). The most reliable way to compare the CDS entries is to extract the drive characters to `curr_cdh` (CHAR) fields of the `cds` table. The phantom drive is simply the string such as "the 3rd H" which appears when you type `DIR NAMI` on a 3rd partition drive. In `MSCDOS`, the equivalent is a string such as "H A".

The special case is to extract `drive_n` (DWORD) for subtraction of 0. The redirector must check

## Another Detection Method: The Network UserVal

*Tim Farley*

The methods used in `is_dev` call for `is_dev()` may not be practical for all redirectors, such as those that own a whole list of drives in the CDS.

Also, comparing the string in the CDS is not practical on some network redirectors, because it might contain one of several server or volume names (and even if you recognized them, that might not conclusively identify the drive as yours). (For instance, suppose you had two different LANs loaded, and you had servers on both LANs that had the same name.)

A different method requires callers to set the `UserVal` parameter in `CX` when calling the `INT 21h` AH:5F03h Make Network Connection function (this shows up at the redirector as an `INT 2Fh` AX:111Eh, with 5F03h on the stack). This `UserVal` is stored at offset 40h in the CDS and is a magic value that identifies your drives. For example:

```

if (cds[drive] & flags & NETWORK)
    if (cds[drive] & magic_parameter == my_magic_number)
        return CALL_IS_FOR_ME,

```

NetWare 4.0 appears to use this, as it requires callers to put "NW" (574Eh) in CX when calling INT 21h AX=5F03h. Of course, your redirector could store the magic value at offset 4Dh itself (if appropriate) and not rely on the caller. Comparing this one word is far easier than comparing an entire string.

**Handling a Read** The redirector's function has finally decided that a call is intended for one of its drives. Let's say the call is a Read; in other words, that a program has called INT 21h AH=3Bh (File Read) or INT 21h AH=3Dh (FCB Random Read) and that DOS has booted this driver to an INT 2Fh AX=1108h call, which has wound up on redirector's doorstep. How does Phantom handle the subfunction 8 Read call?

Looking back at redirector.c in Listing 8-28, you can see that, given a supported redirector subfunction number in AH, this subfunction is called through a dispatch table.

```
if (fixmap[AH] != unsupported)
    dispatch_table[AH]();
```

As you can expect, dispatch table is nothing more than an array of function pointers.

```
PROC dispatch_table[] {
    inquiry, /* 0x00h */
    rd, /* 0x01h */
    unsupported, /* 0x02h */
    wd, /* 0x03h */
    unsupported, /* 0x04h */
    ed, /* 0x05h */
    clfill, /* 0x06h */
    cmetfil, /* 0x07h */
    readfil, /* 0x08h */
    writfil, /* 0x09h */
    //
};
```

Having said that, short story long, the punchline is that when DOS calls INT 2Fh AX=1108h, the Phantom redirector ends up calling the readfil function shown in Listing 8-30. Thus we can start to deal with the standard INT 2Fh AX=1108h trace we produced earlier.

```
21/3f: Read 0005
2f/1108
    redirector() (listing 8-28)
        is_call_for_us() (listing 8-29)
        readfil() (listing 8-30; see alternate in listing 8-36)
21/3f: Read done 0005
```

### Listing 8-30: PHANTOM.C readfil() (INT 2Fh AX=1108h)

```
/* DOS System File Table entry. all DOS versions. NOTE!!! This is
 * slightly different from the standard DOS SFT structure. Some
 * of the fields are for the redirector's use to treat as
 * it sees fit. Others, of course, are maintained and used by
 * DOS, so the * meaning must never be changed. */
typedef struct {
    uint handle_count, open_mode,
    uchar file_attr,
    uint dev_info_word;
    uchar far * dev_drvr_ptr;
    uint start_sector, file_time, file_date,
    file_size, file_pos;
    uint re_sector, abs_sector, dir_sector, dir_entry_no,
    char file_name[11];
    ulong share_prev_sft;
    uint share_net_machine_num, owner_psp;
```

```

// other fields
) SFTREC, far* SFTREC_PTR,
void fat (uint err) { r.flags |= FCARRY, r.ax = err, }
void read_data(long far *file_pos_ptr, uint *len_ptr, uchar far *buf,
  uint start_sector, uint far* last_rel_ptr, uint far* last_abs_ptr),
// Read from file - subfunction 0Bh
void readfil(void)
{
  SFTREC_PTR p = (SFTREC_PTR) MK_FP(r.es, r.di),
  if (!p || !open mode & 1) { fail(5), return, } // access denied

  if ((p->file_pos + r.cx) > p->file_size)
    r.cx = (uint) (p->file_size - p->file_pos),
  if (!r.cx) return; // nothing to do
  // fill caller's buffer and update the SFT for the file
  read_data(&p->file_pos, &r.cx, ((VS_SDA_PTR) sda_ptr)->current_dta,
    p->start_sector, &p->rel_sector, &p->abs_sector);
}

```

Know that way, just. That's all reading files involves. When presented this way, it looks like you could write a DOS disk, the world's most valuable piece of code, one weekend while your spouse is on vacation.

Now, if you're writing a DOS file system, it's a bit harder than this. For example, later on, Listing 8-14 will show a better version of `readfil` with critical error handling. As a RAM disk that handles AMI rather than a sector or two, it's a bit more complex. The bottom line is that it's a bit more complex.

You can see that `readfil` takes the caller's SFT and treats it as an SFTREC\_PTR. This matches what is called about the Read subfunction in the redirector specification later in this chapter.

#### Subfunction 0Bh

##### Read from file

Inputs	ES:DI → SFT for file to read from CX = count of bytes to read SDA_CURR_DTA → user buffer to read data into
Outputs	Carry set = error code in AX if error encountered if no error, CX = bytes actually read SFT updated

In fact, the `Platonic readfil` function is little more than a C version of this specification, taking parameters as the SFT and updating the SFT. The `Read` function reads at the current file position stored in the SFT with the only other parameters needed: the number of bytes to read (CX) and the address of the caller's buffer (DI). All the updating of the SFT is done in the lower-level `read_data` function which, however, knows nothing about SFTs. The function `readfil` passes read data—the addresses of all the SFT fields that need updating.

Back in HELSA, Listing 8-19, we dumped out some SFT fields. However, a key value not shown in HELSA is the current file position. This is maintained not only in the SFT file\_pos field but also in the sector and abs\_sector. `Platonic` treats part of the SFT in its own way, as noted earlier, a redirector can contribute to the meaning of some fields in the SFT. These fields are used and updated by `readfil` and of course also by `writfil` (not shown).

These fields are also updated when a program calls INT 21h AH=42h Move File Pointer a/k/a `beck`. Interestingly, for DOS 4.0 and higher, these `beck` calls usually don't appear at the redirector's doors, but instead, DOS handles with the SFT directly without bothering the redirector. There's one caveat, though, to do with seeking to the end of a file, in which DOS will call the redirector. See the disassembly of `beck` in Chapter 6, Listing 6-20.



**The Phantom XMS File System** At this point, we have satisfied the redirector specification. The internals of the read\_data() function don't matter for the purposes of understanding the redirector. For a particular read\_data() could produce random data such as stock market prices on demand.

However, let's keep going and see how Phantom organizes the file system. Because Phantom implements a FAT-based file system, it is worth examining. It only reinforces the points made much earlier in this chapter about the FAT and DOS directory structures. The read\_data() function is shown in Listing B-31.

The file system implemented by Phantom uses XMS extended memory. The way that it organizes the storage of directories and file information and data is similar, but not identical to the DOS FAT-based file system. However, there is one requirement that a redirector use a FAT file system—that in fact is a little reason to use the redirector in the first place. Since it treats XMS memory as though it were simply a hard disk formatted into sectors, it would be relatively easy to modify Phantom to work on a real, as opposed to virtual, disk. There are separate interesting issues involved with sequential access media, such as tape.

The Phantom file system is sectorized, but into 1024-byte chunks of XMS instead of 512-byte chunks of disk. Phantom divides its disk space XMS allocation into two areas:

- The FAT only occupies. Since a FAT entry that is a WORD in size allows for approximately 64k clusters, a 1024-byte cluster allows the FAT to manage a full 64-MB of XMS memory. Therefore, this portion removes the distinction between sectors and clusters and deals exclusively in sectors.
- The file system area, which accommodates not only the file and subdirectory data and the root directory as well, which Phantom treats as simply another subdirectory. Recall that DOS provides special treatment for the root directory, assigning it a fixed contiguous size not accessible through the FAT.

Apart from this, Phantom emulates a standard DOS file system fairly closely. It uses the same directory entry structure and the same FAT management techniques as DOS. It is worth stressing again that none of this similarity is required of a redirector.

#### Listing B-31: PHANTOM C read\_data()

```
#define SECTOR_SIZE 1024 /* 1024b/sector allows for 64k of XMS */
uchar sector_buffer[SECTOR_SIZE], /* general purpose sector buffer */
      *last_sector = 0xffff, /* last sector read into sector buffer */
/* Uses FAT to find next sector in chain for current file/directory */
extern uint next_fat_sector(uint abs_sector),
void read_data(long far *file_pos_ptr, uint *len_ptr, uchar far *buf,
               uint start_sector, uint far* last_re_ptr, uint far* last_abs_ptr)
{
    uint start, rel_sector, abs_sector;
    uint count, en = *len_ptr;
    start = (uint) (*file_pos_ptr / SECTOR_SIZE);
    if (start < *last_re_ptr)
    {
        rel_sector = 0;
        if ((abs_sector = start_sector) == 0xffff) // end of FAT chain
        {
            *len_ptr = 0;
            return;
        }
    }
    else
    {
```

```

    rel_sector = *last_rel_ptr,
    abs_sector = *last_abs_ptr,
}

while ((len)
{
    start = (uint) (*file_pos_ptr / SECTOR_SIZE);
    if (start > rel_sector)
    {
        if ((abs_sector = next_FAT_sector(abs_sector)) == 0xFFFF)
        {
            *len_ptr -= len; goto update_sectors;
        }
        rel_sector++,
        continue;
    }
    i = (int) (*file_pos_ptr % SECTOR_SIZE),
    count = min((uint) SECTOR_SIZE - i, len);
    if (count < SECTOR_SIZE)
    {
        if (! get_sector(abs_sector, &sector_buffer))
        {
            *len_ptr -= len; goto update_sectors;
        }
        last_sector = abs_sector,
        memcpy(buf, (uchar *) &sector_buffer[i], count);
    }
    else if (! get_sector(abs_sector, buf))
    {
        *len_ptr -= len; goto update_sectors;
    }
    len -= count;
    *file_pos_ptr += count;
    buf += count;
}

update_sectors:
    *last_rel_ptr = rel_sector,
    *last_abs_ptr = abs_sector,
}

```

The read data function contains a standard loop over the number of bytes of data the caller requested. In this FAT-like system, there's effectively one sector per cluster, each 1024 bytes. If the number of bytes requested exceeds one sector cluster read data has to use the FAT to find the location of the next sector, etc. The function calls next\_FAT\_sector (see PHANTOM.C on disk), which uses the sector as an index into the FAT. Of course, on FAT systems that have to read on a FAT sector from a VMS disk. This operation involves first writing the current FAT sector back out to disk.

As happens with a real disk, Phantom reads entire sectors at a time. If the remaining bytes the caller requests are fewer than a full sector read data can't read directly into the user's buffer, so it uses sector buffer, then the bytes copied into the user's buffer.

Since Phantom operates on the host, it does not implement a buffer pool, but most redirectors should. We can get from redirector drives is not by default integrated into the DOS buffers pool, a redirector could call other DOS functions such as INT 21h AX=1210h Find Unreferenced Disk Buffer, or INT 21h AX=1201h Make Buffer Most Recently Used. Note also that data from redirector drives is not cached by programs such as SmartDrive, which work off INT 13h. This is especially important for slow media such as CD-ROM. MSCDEX incorporates its own sector buffers (the location and number of which you can set with MSCDEX.msc) but there is now a market for third-party CD-ROM accelerators (or caches). As noted earlier, SmartDrive 4.2 and higher, included with MS-DOS 6.2, can cache data from CD-ROM drives.

In Listing 8-32, `read_data()` gets the actual data by calling `get_sector()`. This is a macro in `PHANTOM.C`:

```
#define get_sector(sec, buff)
    xms_copy_to_real(xms_handle, (ulong) SECTOR_SIZE * (sec), \
        SECTOR_SIZE, (uchar far *) (buff))
```

The `xms_copy_to_real()` function is in turn just a C wrapper around XMS function `0Bh` (Move Extended Memory Block). Thus, our trace of a DCRS read for a Phantom drive now looks like this:

```
21:3F: Read 0005
2f:1108
    redirector() (listing 8-28)
    io_call_for_usr() (listing 8-29)
    readf(7) (listing 8-30, see alternate in listing 8-38)
    read_data() (listing 8-31)
        XMS function 0Bh
21:3F: Read done 0005
```

In other words, our initial DCRS read turns into a call to XMS function `0Bh`—exactly as one would hope from an XMS RAM disk!

**Handling an Open** Having seen how Phantom handles a DCRS read, we'll now step back to consider how it gets opened in the first place. Rather than the Fat-based Open Create call, we'll look at the `openp()` Open File call (INT 21h AH 3Dh), which DCRS turns into a redirector Open (INT 21h AX 11h). The byte pointer handler for this subfunction, `opnfi`, is dispatched via listing 8-28 in the same way as the `io_opnfi` function is shown in listing 8-32.

In Listing 8-32, `fc`, `ptr`, and `srch_ptr` are, among several Phantom global far pointer variables, pointers to entries in DOS's file SDA and other parts of DCRS. This sounds disgusting, but the `fc` table entry's `instat` depends on the redirector giving direct access to DOS's own global `fc`. An important note: names such as `fcname_ptr` have nothing to do with FC files and merely refer to 11-character filenames.

### Listing 8-32: PHANTOM.C `opnfi()` (INT 21h AX=1116h)

```
/* these are version-independent pointers to various frequently used
   locations within the various DOS structures */
// note: fcname_ptr has nothing to do with FCBs! just 11-char filename
char far * fcname_ptr, // ptr to 1st FCB-style name in SDA */
uchar far * srch_attr_ptr; // ptr to search attribute in SDA */

extern int contains_wildcards(char far* path);
// ffirst(), fill_sft(), init_sft(); see below

void opnfi(void)
{
    // DOS ca is the redirector with (S-D) pointing
    // to an uninitialized SFT entry.
    SFTREC_PTR p = (SFTREC_PTR) PK_FP(r.es, r.di);

    if (contains_wildcards(fcname_ptr)) { fail(3), return; }

    // opening a file requires first finding the file
    *srch_attr_ptr = 0x27, // Archive+System+Hidden+ReadOnly+Rmwa;
    ffirst(
    if ( r.ax )
    {
        // and then using the ffirst search record to fill in the SFT
        fill_sft(p, TRUE, FALSE);
        in_sft_ptr,
    }
}
```

file read. `openfile` is a real find. It does little more than call `find` (find first) and then fill `str` and `direc_ptr`. `findnext` (`Open + get`) manipulates `SEI` from `DCS + find first + set SEI`.

The first time it is seen is Listing 8-33. Listing 8-33 is the same function that would be called if a program did an `IN + 2th AH 4th` `find first` on a sector-by-sector. Before calling `first`, `openfile` sets the sector attribute in the SDA, such as `attr_ptr = 0x20`, indicating "for those readers who do not care to see DOS programming's references around at disk heads that systems hidden and read on files." The fourth in addition to normal files. This matches DOS behavior on local drives, where you can't get open a hidden file although you can't DIR find first one.

### Listing 8-33: PHANTOM.C `ffirst()` (INT 2fh AX=111bh) and `fnext()` (AX=111Ch)

```
char far * filename_ptr,          /* ptr to 1st filename area in SDA */
SRCHREC_PTR srchrec_ptr,        /* ptr to 1st Search Data Block in SDA */
/* Finds the sector number of the start of the directory entries for
the supplied path */
extern int get_dir_start_sector(char far* path, uint far* abs_sector_ptr),
void ffirst(void)                /* findfirst - subfunction 1Bh */
{
char far* path;
int success;

/* Special case for volume-label-only search must be in root */
if (path == (char*)srchrec_ptr == 0x00)
filename_ptr = _fstrchr(filename_ptr, '\\');
*path = 0;

if (path) *path = '\\';
success = get_dir_start_sector(filename_ptr, &srchrec_ptr->dir_sector);
if (! success) { fail(3); return; }

memcpy(&srchrec_ptr->srch_mask, (char*)name_ptr, 11);
srchrec_ptr->dir_entry_no = -1;
srchrec_ptr->attr_mask = *srchrec_ptr;
srchrec_ptr->drive_no = (uchar)four; /* drive no | 0x80 */

/* ffirst() is embedded call to fnext() admittedly looks a little
odd. This arises from the view that findfirst is simply a
findnext with some initialization overhead. findfirst has to
locate the directory in which findnext is to iterate, and
initialize the SDB state to 'point to' the first entry. It
then gets that first entry, using findnext. ffirst() does
initialization, and fnext() is the "workhorse." */
fnext();

/* To mimic DOS behavior, a findfirst that finds no matching entry
returns an error 2 (file not found), whereas a subsequent findnext
that finds no matching entry should return error 18 (no more
files). Having just done the embedded fnext(), we need to turn any
error return value into what is appropriate for ffirst() */
if (eax == 18) eax = 2;
}

DIRREC_PTR dirrec_ptr,          /* ptr to 1st found dir entry area in SDA */
/* Get next directory entry that matches specified mask, continuing
from the supplied starting position (from the previous find) */
extern int find_next_entry(char far* mask, uchar attr_mask,
char far* filename, uchar far* attr_ptr, ulong far* file_time_ptr,
uint far* start_sec_ptr, long far* file_size_ptr,
uint far* dir_sector_ptr, int far* dir_entryno_ptr);
void fnext(void)                /* findnext - subfunction 1Ch */
{
if (! find_next_entry(srchrec_ptr->srch_mask,
srchrec_ptr->attr_mask, dirrec_ptr->file_name,
```

```

    &dirrec_ptr->file_attr, &dirrec_ptr->file_time,
    &dirrec_ptr->start_sector, &dirrec_ptr->file_size,
    &srchrec_ptr->dir_sector, &srchrec_ptr->dir_entry_no)
    fat16_t);
}

```

Aside from some special handling for volume labels, Find First calls `set_dir_start_sector` (not shown, see P.15N10M1) on disk to locate the requested file. If the file is located, `first` fills in the first Search Data Block in the SDB.

We're looking at `first` because, as you saw in Listing 8-52, the `openfcb` function in `phantom` internally calls `first` to locate files. Having called `first` to locate the requested file to open, `openfcb` proceeds to "be actual open" which involves doing more than setting the correct fields in the `fsentry` of SFI entry that DOS passes as FSFI. `phantom`'s `openfcb` calls `init_sft` and `init_sft` to initialize the SFI using the SDB structure pointer filled in by `first` and the directory entry filled in by `first` (see the comments in Listing 8-53 to see why `first` calls `init_sft`). The `init_sft` and `init_sft` functions are shown in Listing 8-54.

### Listing 8-54: PHANTOM.C FILE `sft()` and `init_sft()`

```

#define FREE_SECTOR_CHAIN(sec) \
    while (sect != 0xffff) (sec) = set_next_sector((sec), 0)

extern ulong dos_fctime(void); // call 2F/120D

void init_sft(SFTRC_PTR p, int use_found_T, int truncate)
{
    fmemcpy(p->file_name, fcbname_ptr, 11);
    if (use_found_T)
    {
        p->file_attr = dirrec_ptr->file_attr,
        if (truncate)
        {
            FREE_SECTOR_CHAIN(dirrec_ptr->start_sector);
            p->start_sector = 0xffff;
            p->file_time = dos_fctime(), // includes date, calls 2F/120D
            p->file_size = 0L;
        }
    }
    else
    {
        p->start_sector = dirrec_ptr->start_sector,
        p->file_time = dirrec_ptr->file_time,
        p->file_size = dirrec_ptr->file_size,
    }
    p->dir_sector = srchrec_ptr->dir_sector,
    p->dir_entry_no = (uchar) srchrec_ptr->dir_entry_no,
}

else
{
    p->file_attr = (uchar) *stack_param_ptr, /* Attr is top of stack */
    p->file_time = dos_fctime();
    p->start_sector = 0xffff;
    p->file_size = 0;
    p->dir_sector = srchrec_ptr->dir_sector;
    p->dir_entry_no = 0xff;
}
}

extern void set_sft_owner(SFTRC_PTR sft), // call 2F/120C
void init_sft(SFTRC_PTR p)
{
    /* Initialize the supplied SFI entry. Note the modifications to the open
mode word in the SFI. If bit 15 is set when we receive it, it is an
FCB open, and requires the Set SFI Owner internal DOS function to be

```

```

ca. ed we don't understand this, but this is what MSCDEX does. */
if (p->open_mode & 0x8000)
    p->open_mode |= 0x00FD; // File is being opened via FCB
else
    p->open_mode &= 0x000F;
// Mark file as being on network drive, unwritten to
p->dev_info_word = (uint) (0x8040 | (uint) our_drive_no);
p->file_pos = 0;
p->re_sector = 0xffff;
p->abs_sector = 0xffff;
p->dev_drv_ptr = MskL;

if (p->open_mode & 0x8000) // File is being opened via FCB
    set_sfc_owner(p); // Call 2F/120C
}

```

The `FILE` and `FILE` structures set fields in the SFD from a variety of sources:

- As noted, some SFD fields are set from the first to end directory entry in the SDA (`dirrec_ptr`) and the first Search Data Block in the SDA (`search_ptr`), which are themselves set by `first` and `next`.
- The SFD `ctime` and `ltime` combined into `time` is set using DOS's internal function `INT_21h AX=120Dh` (low number).
- The SFD `drive` field would be set based on our drive no and the magic number `0x8040`, which indicates a non-FAT file system has not been written to bit 0.
- The SFD `open_mode` is field is set with an odd, only semi-understood way for FCBs. Bit 15 in the open mode indicates an FCB Open.
- For an FCB Open, `set_sfc_owner` calls `set_sfc_owner`, which is a C wrapper around `INT_21h AX=120Ch`. Among other things, this DOS internal function uses the current PSP to set the SFD's `owner` field.

For the `FILE` or `FILE` structures, SFD field types show fields such as the above that the redirector sets. For the `FILE` structures, those fields that DOS uses alone for the redirector's own internal use (see DOS FILE STRUCTURE, step 8.24) and those fields that DOS sets.

When `Phantom` opens a path, it sets the result as a field in SFD entry. As you saw in Listing 8-30, many fields are copy operations, such as `p`-file `pos` and `p`-start sector, are subsequently used in `readFile`. Of course, these SFD fields are also changed by other DOS functions such as `Write` and `Seek`.

**Handling Chdir** Having showed at how the `Phantom` handles the `FILE` CD set, we briefly examine `Directory` management. How the `CD` change directory command implemented.

When a program wants to `INT_21h AH=3Bh` Change Directory to a redirector drive (a drive whose CDS has been set up), DOS generates an `INT_21h AX=1105h`. In the `Phantom`, this is handled with a function stream in all its glory in Listing 8-35.

### Listing 8-35: PHANTOM.C cd() (INT 2Fh AX=1105h)

```

char far * filename_ptr, // ptr to 1st filename area in SDA */
char far * current_path, // ptr to current path in CDS */
void cd(void) // Change Directory - subfunction 05h */
{
    /* Special case for root */
    if ((*(filename_ptr) != '\\') || (*(filename_ptr) == 1))
    {
        /* can't make directory with * or ? in name */
        if (contains_wildcards(filename_ptr)) { fa=(3); return; }
        *srch_attr_ptr = 0x10; // look for directory
        ffilesC(); // calls fnext(), which sets dirrec_ptr
    }
}

```

```

    if (r_ax == 0) { if (dirrec_ptr->file_attr & 0x10) { fail(3); return; }
    }
    fstrcpy(current_path, filename_ptr);
}

```

Well, guess it didn't do anything but call `first` to verify that the specified directory string is valid and to place the string into the current CDS. This makes sense. If you were to use a debugger to look at the CDS in memory, it would change your current directory, so of course the DOS `Change` (directory) function does little more than edit the CDS.

**Handling Mkdir** Just now, no directories come into existence in the first place, so that the call to `first` in Listing 8-35 can find them. DOS directories are, of course, created with the `MD` command. Make directory. This command calls `INT 21h AH 39h Mkdir`, which returns for a redirector base call `INT 21h AX 1103h` (you saw this in an earlier code fragment from a disassembly of DOS). In the Phantom, this is handled by the `md` function shown in Listing 8-36.

### Listing 8-36 PHANTOM C `md()` (`INT 2Fh AX=1103h`)

```

#define put_sector(sec, buf) \
    xms_copy_from_real(xms_handle, (ulong) SECTOR_SIZE * (sec), \
    SECTOR_SIZE, (uchar far *) (buf))

void md(void) /* Make Directory - subfunction 03h */
{
    /* special case for root */
    if (!filename_ptr[0]) { if (!filename_ptr[1]) { fail(5); return; }
    }
    if (!contains_wildcards(filename_ptr)) { fail(3); return; }
    *arch_attr_ptr = 0x3f, /* 0x3f = everything
    ffirst(); // if ffirst() succeeds,
    if (r_ax == 0) { fail(5); return; } // directory already exists'
    if (r_ax == 2) return; // we WANT error 2 (not found)
    // if any component part of path is wrong, we'll return error 3
    /* Although we initialize a directory sector, we actually don't need to,
    since we do not create . or .. entries. This is because
    a redirector never receives requests for . or .. in Cdir - DOS
    resolve the absolute path before we get it. If you want to
    see dots in DIR listings, create directory entries for them after
    put_sector. But then you must take account of them in Rmdir. */
    last_sector = 0xffff;
    memset(sector_buffer, 0, SECTOR_SIZE);
    if (!id_exec_ptr->start_sector.next_free_sector()) { fail(5); return; }
    set_next_sector(dirrec_ptr->start_sector, 0xffff);
    last_sector = dirrec_ptr->start_sector;
    if (!put_sector(dirrec_ptr->start_sector, &sector_buffer))
        ( to 1(5), return; ) // access denied
    /* Finally, create entry for this directory. archrec_ptr and
    dirrec_ptr were set in fnext, called from ffirst. */
    if (!create_dir_entry(&archrec_ptr->dir_sector, NULL, filename_ptr, 0x10,
    dirrec_ptr->start_sector, 0, dos_fsize())) // Listing 8-37
        ( to 1(5), return; ) // access denied
    succeed();
}

```

Well, `MD` is one of its drives. Phantom first has to get your new directory string through a series of tests:

- Does it need to create a root directory (error 5 - access denied)?
- Does the file contain wildcards? (path not found)
- Does the path already exist? (access denied)

At first it may look strange that `md()` calls `first()` and, if `first()` succeeds (`ax == 0`), `md()` fails. But if you're creating a new directory, it cannot already exist, so for `md()` to succeed the `first()` call must fail. Next `md()` looks for something in hopes of finding it. Here, `md()` looks for the directory in the hopes of not finding it. The `md()` function can only proceed if and only if `first()` returns error 2 (file not found).

On the other hand, if there are multiple path components, all the higher-level component parts must exist. For example, for an MSD-FDD:BAR:BAZ to succeed, both FDD and FDD:BAR must already exist. If any don't, `md()` fails with error code 3 (ret. mod. from `first()`). The `md()` function can only proceed if `first()` returns error 2 (file not found).

At this point, `md()` can go ahead and make the directory, which basically involves some sector manipulation. If they were DDOS rather than the Phantom, it would of course manipulate clusters rather than sectors. As seen in `md()`, the steps involve getting a free sector, setting its FAT entry (FAT sector pointer) to 0FFFH, writing the sector to disk, and finally creating the entry for the new directory in its parent directory.

It won't be so much obvious from `md()` how the new entry is plugged into the parent directory, that is how MSD-FDD:BAR would end up smacking an entry for BAR into FDD's directory. The very end of the function, `create_dir_entry()`, to create this directory entry, but it isn't clear how `md()` knows the parent directory entry which is plugging it. As usual in the Phantom, the answer lies in the first function, `get_dir_entry()`, which is `md()`'s job to ensure that the target directory doesn't already exist, but doesn't like the idea. `first()` itself also implicitly locate the new directory's parent.

Case 3: The location of the new directory's parent directory. `md()` creates the new directory entry by calling `create_dir_entry()`. As shown in Listing 8-37, this function walks through a directory, looking for a sector for a deleted entry. If a free entry can't be found in the current directory sector, create\_dir\_entry() searches FAT to find the next directory sector. If there isn't a next directory sector, create\_dir\_entry() calls `get_dir_entry()` to create one. This new directory sector, assuming it can be allocated, is to be linked to the FAT to the current one; this FAT link is made by `set_next_sector()`, also shown in Listing 8-37.

### Listing 8-37 PHANTOM.C create\_dir\_entry() and set\_next\_sector()

```
int create_dir_entry(uint far *dir_sector_ptr,
    uchar far *dir_entryno_ptr, char far* filename, uchar file_attr,
    uint start_sector, long file_size, ulong file_time)
{
    uint next_sector, dir_sector = *dir_sector_ptr,
    DIRREC* dr = (DIRREC*) &sector_buffer,
    int i;

    for (i = 0;
        {
            if (dir_sector != last_sector)
            {
                if (!get_sector(dir_sector, sector_buffer))
                    return FALSE;
                else
                    last_sector = dir_sector;
            }
            for (i = 0; i < DIRREC_PER_SECTOR; i++)
            {
                if (dr[i].file_name[0] && dr[i].file_name[0] != (char) 0xE5)
                    continue; // looking for unused or deleted entry
                memcpy(dr[i].file_name, filename, 11);
                dr[i].file_attr = file_attr;
                dr[i].file_time = file_time; // includes date
                dr[i].file_size = file_size;
                dr[i].start_sector = start_sector;
                *dir_sector_ptr = dir_sector;
            }
        }
    }
}
```



```

    if (dir_entryno_ptr) *dir_entryno_ptr = (uchar) i;
    return put_sector(dir_sector, &sector_buffer);
}

// no free dir entry: get next, or allocate new
if ((next_sector = next_FAT_sector(dir_sector)) == 0xFFFF)
{
    if (! (next_sector = next_free_sector()))
        return FALSE;
    set_next_sector(dir_sector, next_sector);
    set_next_sector(next_sector, 0xFFFF);
}
dir_sector = next_sector;
}
}

uint FAT_page[FATPAGE_SIZE]; /* buffer for FAT entries */
int cur_FAT_page = -1; /* index of FAT page in buffer */
int FAT_page_dirty = FALSE; /* Has current FAT page been updated */
uint free_sectors; /* unallocated sectors on XMS disk */

/* Checks that the page of FAT entries for the supplied sector is in
the buffer. If it isn't, go get it, but write back the currently
buffered page first if it has been updated. */
extern int check_FAT_page(uint abs_sector);

/* Update the FAT entry for this sector to reflect the next sector
in the chain for the current file/directory */
uint set_next_sector(uint abs_sector, uint next_sector)
{
    uint save_sector;
    if (! check_FAT_page(abs_sector)) return 0;

    save_sector = FAT_page[abs_sector - (cur_FAT_page * FATPAGE_SIZE)];
    FAT_page[abs_sector - (cur_FAT_page * FATPAGE_SIZE)] = next_sector;
    if (save_sector > next_sector)
    {
        FAT_page_dirty = TRUE;
        if (! save_sector) free_sectors--;
        else if (! next_sector) free_sectors++;
    }
    return save_sector;
}
}

```

Any time a sector of a file was found or created a free directory entry, it initializes the directory using the file system's parameters. It name, file size, start sector, and sectors. Initially, the directory sector is written back to disk using put\_sector, a macro that calls xms\_copy\_in\_real. As with xms\_copy to real, this is a C wrapper around XMS function 0Bh. More Extended Memory Block. In the Phantom XMS RAM disk "writing" a sector involves a copy from conventional real mode memory to XMS, whereas "reading" a sector involves a copy from XMS to conventional memory.

When the creation of a directory entry in Listing 8-37 is successful, it still helps to examine the code. In this chapter, we used directory entries without giving much thought to how fields such as file's starting sector come to be there in the first place. In create\_dir\_entry and main or get\_dir of Phantom, you can view the same structures from the operating system's viewpoint. As you'd expect, these are basically part of the operating system.

This code is our trace through the Phantom code. Again, complete source code is available on disk as UNDOOR.ZIP\HAPPHPHANTOM.C.

### Differences Between DOS Versions

Phantom works with DOS versions from 3.10, when Microsoft introduced the redirector interface, through to DOS 6.0. The interface has changed little in that time, except in one or two important

areas. Perhaps the most predictable change is that some subfunctions have been added to cater to new functions added to the INT 21h interface.

In DOS 4.0, for example, the Extended Open File function was introduced to make all types of file ops available through one call. At that time, a corresponding new redirector interface subfunction number, 21h, was added to handle Extended Open. The SDA also changed with DOS 4.0, and several existing 8-24 SROP, M-FILES, SPROP, MMODE, and SROP ATTR fields were added to support the special open functionality.

Another subfunction number introduced with DOS 4.0 was 21h. Some of the DOS internal comments use the ~~unsubstantiated~~ DOS function 57h, which appears to trigger 21h at the redirector interface. However, both DOS function 57h and the redirector subfunction 21h are among those DOS 4.0 files which do not appear in late-DOS versions, so we won't worry about them further. The following redirector specification quite deliberately writes off anything specific to DOS 4.0 as a dead end.

### The Network Redirector Specification

The following table presents the known redirector subfunctions, with usage parameters, and notes. Remember that DOS merely defines this specification and that any given redirector must supply the actual functions that meet this specification.

#### Subfunction 00h

##### Installation Check

Inputs: None

Outputs: AL = 00h not installed, OK to install  
AL = 01h not installed, not OK to install  
AL = FFh installed, OK to install

- 1 A redirector should call this subfunction at initialization and should not load if the subfunction returns 01h or FFh. It returns 00h or 11h if it is OK to load. These two values allow a redirector to opt to call, load, or install another redirector, so a ready present. Once installed, the redirector should respond to other redirectors that call this subfunction, usually with 11h or AL, unless there is a reason to disallow subsequent redirectors to load, in which case the redirector should set AL to 01h.
- 2 A redirector should never handle this call. This convention means that it is up to the most recently loaded redirector to decide whether any further redirectors may be loaded. You can see this by taking the INTCHEM program from Chapter 6, run INTCHEM 21:410h, and notice that the most recently loaded redirector takes the call. Again, the whole issue of chaining redirectors is complicated.
- 3 MSCDEX looks for a previous instance of the file, pushes 0BADAh on the stack before calling Sys 21:4A:1400h. After return from the INT 21h, it tests if the top of the stack is still 0BADAh. If not, i.e., the INT 21h handler has changed 0BADAh to something else, it concludes that MSCDEX is already loaded, it exits instead of going ISR. Meanwhile, the MSCDEX handler or subfunction always sets the word at the top of the stack to 0BADAh. Note what David Markin calls the "stack" in this interface: the resident code sets a 0BADAh on the prospective resident code tests only for 00BADAh. At least one programmer in person (MS-DOS 4.01, Networker 4.x, and 5.x) looks for 0BADh on the stack and changes it, not only if there was a generic MSCDEX that loaded before 4.01 Networker. Networker does not change 0BADh to 0BADh, but instead means incrementally, allowing a call to detect that something special is going on.

#### Subfunction 01h

##### Remove Directory

Inputs: 50A.FH1 = fully qualified directory name

Outputs: Carry set, error code in AX if error encountered

- The redirector should either compare the supplied CDS\_PTR pointer in the SDA with the address of the CDS for one of its drives or, preferably, the referenced CDS contents with the contents of its own CDS to determine if the call is intended for its drive.
- The redirector should ensure that it doesn't remove the current directory by comparing SDA\_LEN with the current path in the CDS; attempts to remove the current directory should be failed with error code 16.

**Subfunction 03h****Make Directory**

Inputs: SDA\_PTR = fully qualified directory name  
 Outputs: Carry set, error code in AX if error encountered

See note 1 for subfunction 01h.

**Subfunction 05h****Change Current Directory**

Inputs: SDA\_PTR = fully qualified directory name  
 Outputs: Carry set, error code in AX if error encountered

- See note 1 for subfunction 01h.

- The redirector must update the C\_CUR\_PATH field of the CDS for the drive.

**Subfunction 06h****Close File**

Inputs: ES:DI -> SFT for file to close  
 Outputs: Carry set, error code in AX if error encountered  
 SFT completed if no error

- The redirector should use the bottom 6 bits of the DEV\_INID field of the SFT pointer to by INID to determine whether the call refers to a file on one of its drives (i.e. SFT\_0\_A1-B) and subtract it. There may be an issue involving the DPF field within the SFT.
- Increment decrement C\_HANDLES first field in the SFT and create or update directory information for the file if it was opened for writing (bit 0 of I set). Calling to key into this SFT handle count creates orphaned files. The redirector can decrement the handle count directly (no handle count) or by calling INT 21h AX=1200h.

**Subfunction 07h****Commit File**

Inputs: ES:DI -> SFT for file to commit (flush buffers)  
 Outputs: Carry set, error code in AX if error encountered

See note 1 for subfunction 0.

**Subfunction 08h****Read from File**

Inputs: ES:DI -> SFT for file to read from  
 CX = count of bytes to read  
 SDA\_CURR\_DTA -> user buffer to read data into  
 Outputs: Carry set, error code in AX if error encountered  
 if no error, CX = bytes actually read  
 CX = 0 to indicate end of file  
 SFT updated

- See note 1 for subfunction 0.
- The redirector should also update the F\_POS field in the SFT.
- Don't forget to set CX = 0 to indicate EOF.

**Subfunction 09h****Write to File**

Inputs: ES:DI -> SFT for file to write to  
 CX = count of bytes to write  
 SPA (CURR\_DTA -> user buffer from which to write data

Outputs: Carry set, error code in AX if error encountered  
 if no error, CX = bytes actually written  
 SFT updated

- 1 See note 1 for subfunction 0h
- 2 The redirector should also update the FILEPOS and FILESIZE fields in the SFT
- 3 The call should fail with access denied (5) if the file was opened for reading only (bits 0 and 1 both zero)
- 4 If CX is 0, truncate the file to the current file position

#### Subfunction 0Ah

Lock/Unlock Region of File

Inputs: ES:DI -> SFT for file  
 CX:DX = region offset (DOS 3.0.X)  
 SI = high word of region size (DOS 3.X)  
 Word at top of stack = low word of region size (DOS 3.X)  
 BL = 0 (Lock) or 1 (Unlock) (DOS 6+)  
 DS:DX = LOCKREC for region to lock (DOS 4.0+)

Outputs: Carry set, error code in AX if error encountered

- 1 See note 1 for subfunction 0h
- 2 The redirector is expected to resolve lock conflicts. Loading SHARE has no effect on a redirector. The redirector must do all arbitration itself. This is obvious in the case of a remote work redirector which must keep a state of the server.
- 3 This function only provides locking in DOS 3.1-3.3 with parameters in registers and on the stack. Both locking and unlocking are achieved through this subfunction in DOS 4.0 and above, with parameters in the LOCKREC structure.

#### Subfunction 0Bh

Unlock Region of File

Inputs: ES:DI -> SFT for file  
 CX:DX = region offset  
 SI = high word of region size  
 Word at top of stack = low word of region size

Outputs: Carry set, error code in AX if error encountered

- 1 See note 1 for subfunction 0h
- 2 The redirector is expected to resolve lock conflicts.
- 3 This subfunction is only called in DOS 3.1-3.3 and is superseded by subfunction 0Ah, BL=1 in DOS 4.0 and higher.

#### Subfunction 0Ch

Get Disk Space

Inputs: ES:DI -> CBS for drive

Outputs: AL = Sectors per cluster  
 BX = Total clusters  
 CX = Bytes per sector  
 DX = Number of available clusters

- 1 The redirector should extract the supplied CBS pointer in ES:DI (or in the SDA) as with subfunction 0Bh, with the address of the CBS for its drive or, preferably, the CBS contents with the contents of its own CBS to determine if the call is intended for its drives.
- 2 The units of sectors and clusters are DOS arbitrary and may not be appropriate to the redirector's underlying storage. It is sufficient to return values such that (AL \* CX \* BX)

reflects the size in bytes of the drive as the redirector wants it reported, and that

AL \* C \* V \* D \* Y reflects the amount of free space in bytes that the redirector wants reports to be available. The register usage here limits the size of redirector drives to 1,024 gigabytes.

#### Subfunction 0Eh

##### Set File Attributes

**Inputs:** SDA.FMT = Fully qualified filename  
 SDA.CURR\_CDS → CDS for drive with file  
 Word at top of stack = New file attributes

**Outputs:** Carry set, error code in AX if error encountered

See note for subfunction 01h

#### Subfunction 0Fh

##### Get File Attributes

**Inputs:** SDA.FMT = Fully qualified filename  
 SDA.CURR\_CDS → CDS for drive with file

**Outputs:** Carry set, error code in AX if error encountered  
 If no error, AX = file attributes  
 BX:DI = file size

1 See note for subfunction 01h

2 FILE\_SIZE, FILE\_SIZE, Get File Size depends on the BY:DI file size return value

#### Subfunction 11h

##### Rename File

**Inputs:** SDA.FMT1 = Current fully qualified filespec  
 SDA.FMT2 = New fully qualified filename  
 SDA.CURR\_CDS → CDS for drive with file

**Outputs:** Carry set, error code in AX if error encountered

1 See note for subfunction 01h

2 The redirector can use the SRC FILE, ECH\_ND\_FILE, RIN\_SRC\_FILE, and RNT\_FILE fields of the SDA as a workspace for iterating over source and target filespecs.

#### Subfunction 13h

##### Delete File

**Inputs:** SDA.FMT = Fully qualified filespec  
 SDA.CURR\_CDS → CDS for drive with file

**Outputs:** Carry set, error code in AX if error encountered

1 See note for subfunction 01h

2 The redirector can use the SRC FILE and ECH\_ND\_FILE fields of the SDA as workspace for iterating over source and target filespecs.

#### Subfunction 16h

##### Open Existing File

**Inputs:** SDA.FMT = Fully qualified filename  
 SDA.OPEN\_MODE = Open mode for file  
 SDA.CURR\_CDS → CDS for drive with file  
 ES:DI → Uninitialized SFT for the file

**Outputs:** Carry set, error code in AX if error encountered  
 SFT completed if no error

1 See note for subfunction 01h

2 The redirector should not set the C\_HANDLES field in the SEI\_DOS mantissa 1 as field result.

3 The type of the open mode will be set for an H B open.

Subfunction 17h  
Create/Truncate File  
Inputs: SOA.FM1 = Fully qualified filename  
ES:DI => Uninitialized SFI for the file  
SOA.CURR\_CDS -> CDS for drive with file  
Word at top of stack = File attribute for file  
Carry set, error code in AX if error encountered  
SFI completed if no error

- 1 See note for subfunction 01h
- 2 INT 21 - ALL SBI Create New File calls this subfunction. For 50h to fail if the file exists, every system with INT 21 must be able to communicate the pre-existence of a file. It does in this system by changing the value at the top of the stack.
- 3 The redirector would not set the C.HANDLE field in the SET\_DOS maintains this field itself.

Subfunction 18h  
Create File without CDS  
Input and Output: Identical to subfunction 17h (see above)?

C: dir \Net 2 > ALL SBI Create New File with a UNC filename such as "0:FOO\BAR" (mg gets FDIS -> no subfunction 18h). It is similar to subfunction 17h, but is only called when the current CDS pointer in SDA has an offset of 0FFFFh (see also subfunction 19h). This call may have been used to support LAN Manager named pipes.

Subfunction 19h  
Find First without CDS  
Input and Output: Identical to subfunction 18h (see below)

Like subfunction 18h, this allows redirectors to manage virtual files on drives without a CDS. FDIS calls this function in its Find First code when the current CDS has an offset of 0FFFFh.

```

FDCB 6C22  C43EA205      LES DI,[E05A2]    ; current CDS
FDCB 6C26  85FFFF      CMP DI, 01       ; offset 0FFFFh
FDCB 6C29  75D6        JNZ DO_2F_111B   ; do regular redir call
FDCB 6C2B  B81911      MOV AX,1119     ; do no-CDS redir call
FDCB 6C2F  CD2F        INT 2F
FDCB 6C30  C3          RET
                DO_2F_111B
FDCB 6C31  261745430080    TEST Word Ptr ES [DI+43],8000 ; CDS network bit
FDCB 6C37  74D6        JZ no_redir
FDCB 6C39  B81B11      MOV AX,111B
FDCB 6C3C  CD2F        INT 2F
FDCB 6C3E  C3          RET
                no_redir

```

This call (and thus the CDS offset 0FFFFh) and to be can be triggered with a command such as "DIR FOO\BAR".

## UNC Filespecs

Tim Farley

**When a redirector is installed that supports the non-CDS calls, you can use UNC-style filespecs directly in DOS calls.** For instance, you could DIR \>SERVER1\VOL1\\*.\* to search the root drive of the volume VOL1 on server SERVER1. Even though it is a front-end hook and not a redirector, the Novell NetWare shell (NETX) supports this behavior. You can also

do this under many of the networks which are redirectors, including Artisoft's LANtastic NetWare Lite 1.0; on the other hand, doesn't support a DIR of a UNC filespec.

The non CDS calls can help if you know the exact network name of a file you need to access but don't want to set up a CDS entry for a fake DOS drive on which to access that file. Some of Novell's own utilities such as LOGIN.EXE use this trick.

Understanding UNC filespecs is crucial to writing a proper network redirector, especially when it comes to implementing the INT 21h AH=5Bh and AH=5Fh calls. Microsoft's *LAN Manager's Programmer's Reference* briefly discusses UNC.

#### Subfunction 0Bh

##### Find First Matching File

**Inputs:** SDA INT = Fully qualified filespec for search  
 SDA SDB = Uninitialized Search Data Block  
 SDA FOUND\_FILE → Directory info buffer for found file  
 SDA SRCH\_ATTR = Search attribute mask for file

**Outputs:** Carry set, error code in AX if error encountered  
 If no error, SDB initialized

See note 1 for subfunction 01h

#### Subfunction 1Ch

##### Find Next Matching File

**Inputs:** SDA SDB = Search Data Block from last Find operation  
 SDA FOUND\_FILE → Directory info buffer for found file

**Outputs:** Carry set, AX = 1Ch if no more files

The redirector should use the bottom 6 bits (mask 0B7) of the DRIVE\_NUM field of the SRC\_HIBLX field of the SDA to determine whether the call continues a previous search on its drive number (0 = A:, 1 = B:, and so forth).

But it, above, is probably wrong. According to David Markin:

The above logic, and the corresponding mask with 3Eh code, will lead to files reads under IBM Lan Server. You may also get a situation where the Phantom will claim a find that a file find request was handled by Lan Server. MSC.DEX does not have this problem because it explicitly tests for 0B400 to be sure it can before claiming the call. Masking with 3Eh will ignore that bit and thus mistake it as any other call. In Business's cooperating redirectors store a 1 based drive in the low 5 bits of the SDB drive\_no field. Lan Server is a cooperating redirector that uses the low 5 bits in the same fashion, thus requiring cooperating redirectors to distinguish themselves with bit 6 set. Cooperating redirectors identify themselves by setting bit 6 in the drive letter. This is the same bit that turns a 1 based drive into a drive letter A-Z.

#### Subfunction 10h

##### Close all Files for Process

In implementing this function, the redirector must maintain a record of all files opened and by which processes on which machines. This is a DOS broadcast called even if no drive in the CDS has the network server. DOS calls this function immediately after using INT 21h AH=31h to close all of a terminating process's open files. The current PSP ptr points to the terminating process.

#### Subfunction 1Eh

##### Do Redirection

**Inputs:** Word at top of stack = Command to execute (e.g., 5F02h)  
 Other inputs depend on command to execute

**Outputs:** Carry set, error code in AX if error encountered  
 Other outputs depend on command to execute

this stack end to INT 21h AH=5Fh. For example, if you want your redirector's drives to show up in INT 21h AX=5F02h, assign hst; see NE1DRA.C Listing 8.15; then you need to implement the next function to determine a call for ax; for the AX=5F03h Make Network Connection call, use the Userfile.NCY; see Another Detection Method. LAN Manager also uses INT 21h AH=5Fh (see *The Dobb's Journal*, April 1993).

#### Subfunction 1fh

##### Printer Setup

**Inputs:** word at top of stack = Command to execute (e.g., 5E02h)  
 Other inputs depend on command to execute  
**Outputs:** Carry set, error code in AX if error encountered  
 Other outputs depend on command to execute

DOS calls this subfunction for INT 21h AH=5Fh, similar to the way that 2F/11E is the back end to 21/5F.

#### Subfunction 20h

##### Flush All Disk Buffers

**Inputs & Outputs:** unknown

This is a DOS1 broadcast, which the Reset Drive function (INT 21h AH=00h) is called, DOS calls INT 21h AX=1120h, even if no CDS has the network bit set.

#### Subfunction 21h

##### Seek From End of File

**Inputs:** ES:DI → SFT for file  
 CX:DX → Offset relative to end of file to position to  
**Outputs:** Carry set, error code in AX if error encountered  
 DI:AX = new file position

This function is almost always called at DOS 4.0 and higher, so don't depend on it to keep you informed of file position changes. The redirector should always use the LPOB field in the SFI to determine the correct file position at which to read or write. As seen in Chapter 6, given an INT 21h AH=42 back, DOS almost always handles directly with the SFI without calling the redirector. Subfunction 21h is always called with a program exec method #2, move from end, and even then only for the FAT Programs when certain SHARE options allow write access. The conditions under which this subfunction is called are so specialized that it appears to have been a special purpose hack.

#### Subfunction 22h

##### Process Termination Hook

**Inputs:** DS = PSP of process about to terminate

Whenever a program exits, DOS issues this subfunction 22h broadcast. The following small INTRSPY script hooks this call to compute a list of programs that have terminated:

```
intercept 2fh function 1fh subfunction 22h
on_entry output (ds 1 B >byte,asc)z,B) (PSP ds 1 exiting)
```

It's much like INTRs that most know when programs exit. Trapping subfunction 22h is far easier than hooking INT 21h, INT 20h, and INT 27h, which is how a front-end hook would wait for process terminations.

#### Subfunction 23h

##### Qualify Path and Filename

**Inputs:** DS:SI → Unqualified filename  
 ES:DI → Buffer for fully qualified filename  
**Outputs:** Carry set, error code in AX if error encountered

DOS appears to support a default name qualifier function that does a very adequate job without support from a redirector. DOS appears to need the assistance of this redirector function only for some some special directory or filename translations. The output of this function (or of the DOS



default routine directly supplies the input for the directory and file subfunctions. This call is a DCS broadcast made via 0 to redirector is running.

**Subfunction 24h**  
Turn Off Remote Printer

Called via DCS 51h kernel subfunction 26h below, returns carry set

**Subfunction 25h**  
**Redirected Printer Mode**  
**Inputs** Word at top of stack = Command to execute (e.g., 5007h)  
Other inputs depend on command to execute  
**Outputs** Carry set, error code in AX if error encountered  
Other outputs depend on command to execute

DCS turns a call to INT 21h AX=5007h Act Redirected Printer Mode AX=51008h Set Redirected Printer Mode, or AX=51009h Flush Redirected Printer Output, into a broadcast of subfunction 25h. Unless an INT 21h handler supports this subfunction, these INT 21h calls perform no action.

**Subfunction 26h**  
**Remote Printer Echo On/Off**  
**Inputs** ES:DI => SFT for file handle 4 (stdin)  
??  
**Outputs:** CF set on error

DCS calls this subfunction when print echoing (P \*Prn% changes state, and stdin has bit 11 ("network spooler") of the device information word set in the NI).

**Subfunction 27h**  
**Remote Copy**  
**Inputs** SI = Source file handle  
DI = Destination file handle  
CX = Bytes to copy (from/to current seek positions)  
BX = NW (574eh, signature for NetWare)  
CF clear if successful, set if failed  
**Outputs** AX = Return code, if error  
05h Access denied: No read rights or no write rights  
06h Invalid handle: One of the file handles is invalid.  
0Bh Invalid format: Improper signature in BX  
11h Device not the same: Both files are not handled by the redirector.  
30h Unexpected network error.  
If successful, current file positions updated for both files

When a particular file seems to be Novell-specific, this function copies one file to another file where both files are on the redirected drive. This can greatly reduce network traffic. This function does not rely on the global from which DCS and can therefore be called directly from an application, or via other file access subfunctions; this function is passed file handles, not SFT pointers.

The Remote Copy subfunction appears to be supported in Netware 4.0; the beta documentation for which provided the basis for the above description. Microsoft LAN Manager has a similar NetRemoteCopy function (INT 21h AX=5134h).

**Subfunction 28h**  
**Extended Open File (DOS 4+)**  
**Inputs:** SDA FN1 = Fully qualified filename  
ES:DI => Uninitialized SFT for the file  
Word at top of stack = File attr for created/truncated file  
SDA SPECOPEN\_ACT = Action codes  
SDA.SPECOPEN\_MODE = Open mode for file  
**Outputs** Carry set, error code in AX if error encountered  
SFT completed if no error  
CX = result code

01h file opened  
 02h file created  
 03h file replaced (truncated)

- 1 See note for subfunction 01h.
- 2 DOS 4.0 introduces this function to provide a unified interface to the functionality supplied by subfunctions 10h and 17h and to support DOS Function 6Ch.
- 3 The redirector should not set the C-HANDLES flag as the SEI-DOS maintains this field itself.

There are other INT 21h API 11h functions that are not part of the network redirector interface. For example, the LAN Manager API for DOS provides asynchronous named pipes using INT 21h API 11h calls. For more information, see Michael Stucky, "The Undocumented LAN Manager and Named Pipe APIs for DOS and Windows," *The Doc's Journal*, April 1993; this appeared in the *DDJ* "Undocumented Corner".

### Using DOS Internal Functions

Because redirectors effectively become part of DOS, they have at their disposal many of the same functions that DOS itself uses internally. As we saw in Chapter 6, DOS uses INT 21h API 12h to support many DOS internal functions. This interrupt-based interface was probably initially designed, or at least thought up, for use by redirectors and presents a wealth of quite useful functions that are only accessible to redirectors and other OS DOS subsystems.

Phantom uses only three DOS internal DOS functions and the only mandatory one is Set SEI Owner (INT 21h API 120Ch). As stated above, one of the benefits of the redirector interface is that it hides within it, except for the existence of the two low-level methods (FCB vs. handle-based), all the DOS programs. The only exception is that, at its opening, the redirector must call Set SEI Owner for FC programs. In 15-seg, in the open-stack word in the SEI passed to the redirector open call (INT 21h API 120h; see above) indicates an FCB open. Failure to call this function in this situation leads to "FCB unavailable" errors.

Phantom also calls Get Date and Time (INT 21h API 1209h) to obtain the system date and time in EBCDIC format. It could equally well convert the BIOS time tick count value for the time, and the MM and YY 1980 fields of the SDA for a date to generate the necessary values for new files by a date-time stamping, but the DOS internal function is right in the spot.

Finally, Phantom calls Check File Characteristics (INT 21h API 1223h) to check file access to S11, C00, and S00, which DOS does on file setup, as before passing control to a redirector. Without this check, Phantom would mistake the get access attributes with names such as S11 and C00.

The following is a list of the INT 21h API 12h subfunctions that a foreign file system redirector might use. The entries do not describe usage in detail; see the Appendix or the Interrupt List on disk, but instead discuss why a redirector might want or not want to use the function.

- 06h Invoke critical error

This is potentially a very important function for a redirector that relies on a physical, network, or other possibly asynchronous medium. It triggers the mechanism that leads to invocation of the documented Int 24h critical error handler. There are two requirements for its proper use:

- Set the appropriate error class, i.e., a suggested action. These are, happily, accessible in the ERR\_CODE, ERR\_CODE\_PTR, W\_CODE, and ERR\_CLASS fields in the SDA.
- Create a dummy-level driver header for the drive and record its address in the DEVDRIVER\_PTR field of the SDA. This is required because DOS uses the device attributes word of the device driver associated with the error to decode the type of error message to display.

Errors generated during open, read, write, close and findfirst, findnext processing are candidates for critical error processing. Because Phantom does not use an unreliable medium, in general a rash assumption about VMS, especially given well-known bugs in the VMS `tcsize` function, it has no need for critical error capabilities.

However, the code fragment in Listing 8-38 shows modifications that might be made to `readfil` if PHANTOM file failure at read time were possible. This code fragment assumes that `read_fil` now returns an int indicating success (non-zero) or failure (zero). This code is an alternate to the version of `readfil` shown earlier in Listing 8-30.

### Listing 8-38: Alternate Version of `readfil()` to Handle Critical Errors

```
// Possible returns from critical_error()
#define CRITERR_IGNORE 0
#define CRITERR_RETRY 1
#define CRITERR_ABORT 2 //ABORT returns control to DOS, not us
#define CRITERR_FAIL 3

uint sp_before_switch; // save area for our sp
// This is the dummy device driver header to keep DOS happy at
// critical error invocation. In this case, the critical error message
// will be 'Error reading drive 0'. Change the dev_attr field value
// from 0 to 0x8000 to generate 'Error reading from PHANTON'
structure {
    long next_hdr;
    uint dev_attr;
    uint strat_entry, intr_entry;
    char dev_name[8];
} dummy_devhdr = {-1,0,0,0,{'P','H','A','N','T','O','N',' '}};

//Set DOS extended error info.
void set_dos_extended_err(uint ax_val, uchar bl_val, uchar bh_val,
    uchar ch_val)
{
    // ax_val -- DOS extended error code
    // bh_val -- error class
    // bl_val -- suggested action
    // ch_val -- locus of error
    ((V3_SDA_PTR) sda_ptr)->err_code = ax_val,
    ((V3_SDA_PTR) sda_ptr)->err_action = bl_val,
    ((V3_SDA_PTR) sda_ptr)->err_class = bh_val,
    ((V3_SDA_PTR) sda_ptr)->err_locus = ch_val,
    ((V3_SDA_PTR) sda_ptr)->devdrv_ptr = (void far *) &dummy_devhdr;
}

//Invoke DOS critical error handler and get user
// return code (Ignore, Retry, Abort, Fail).
int critical_error(uchar ah_val)
{
    int ret;
    //Invoke critical error.
    asm {
        mov ah, ah_val
        push bp
        push si
        push di
        mov di, 000bh // Read error
        mov bp, ds // BP:SI must point to our fake
        mov si, OFFSET dummy_devhdr // device header
        mov sp_before_switch, sp //Save current stack pointer
        cll
        mov ss, dos_ss //Establish DOS's stack, which was
        mov sp, dos_sp // current when we got called.
        sti
    }
}
```

```

mov al, our_drive_no      //Driveletter
push ax                   //Int 24h AX value goes on the stack
mov ax, 1206h             //Invoke critical error.
int 2Fh
mov bx, ds                 //Restore our stack segment (same as DS)
cli
mov ss, bx
mov sp, sp_before_switch // and stack pointer (which we saved).
sti
pop di
pop si
pop bp
xor ah, ah
mov ret, ax
}
return ret;
}

void readf(void)          // Read from File - subfunction 08h
{
SFIREC_PTR p = (SFIREC_PTR) PK_FP(r.es, r.di);
if (p->open_mode & 1) { fail(5); return; }

if ((p->file_pos + r.cx) > p->file_size)
r.cx = (uint) (p->file_size - p->file_pos);
if (! r.cx) return;

/* Fill caller's buffer and update the SFI for the file */
while (read_data(sp->file_pos, &r.cx,
((V5_SDA_PTR) sda_ptr)->current_dta,
p->start_sector, sp->rel_sector, sp->abs_sector) == 0)
{
set dos_extended_errf 0x001E, // Error - Read fault
0x04, // Action - Orderly abort
0x0B, // Class - Media error
0x05, // Locus - Memory

// call the user's critical error handler. This is interesting
// because we are used to seeing critical error from the app's
// point of view; here we see it from DOS's point of view.
switch (critical_error(111110b)) // 0 - 0 indicates read
// 1,2 - 11 indicates file error
// 3 - 1 indicates a low fail
// 4 -- 1 indicates a low retry
// 5 - 1 indicates a low ignore
{
case CRITERR_IGNORE: break;
case CRITERR_FAIL: fail(2), return, // not reached
case CRITERR_RETRY: break;
default: break, // impossible...
}
}
}
}

```

If `readf` encounters an error, it first calls `set_extended_errf` to set the appropriate codes into the SDV; it then calls the `critical_error` function with flags to indicate where the error occurred and what responses are to be shown. These bits, together with the extended error class, action, and locus values are documented in the DOS programmer's reference for INT 21h AH-9h Get Extended Error. The return from that function is the user or application's response to the error.

Although `CRITERR_ABORT` is defined, we never receive control back if that is the selected response. DOS fails the application or command immediately after the Int 24h handler returns.

As a final note on critical errors, redirectors should use the INI 2Fh AH 05h interface to expand critical error messages. For example, MSC DEX uses this interface. There is a detailed discussion of this interface and other error handling issues in Chapter 17 of Geoff Chappel's *DDoS Internals*.

- 08h Decrement SFI Reference Count

Redirector subtraction to close file must decrement the SFI handle count to avoid creating orphaned files. INI 2Fh AX-1208h does just this. However, it is much more efficient to decrement the SFI reference count directly rather than to call this function. For example, the following is lifted from the Phantom's `clslib` function:

```
SFTREC PTR p (SFTREC_PTR) WC_FP(r.es, r.di),
if (p->handle_count) /* If handle count not 0, decrement it */
    - p->handle_count,
```

- 0Ah Perform Critical Error Interrupt

This function is similar to the preferable function 06h. The function 0Ah version appears to need a DFB for the redirected drive; in general, redirector drives don't require DFBs.

- 0Bh Signal Sharing Violation To User

For peer DDBS networks and SHARE LM, use this function to signal that a command or application has attempted to open a file previously opened by FCB or with a sharing "deny." Its principal return is the value of the carry flag, indicating whether to retry the operation.

- 0Ch Set SFI Owner (Set FCB Owner)

As I've noted earlier, this is the one essential DDBS internal call for a redirector. While DDBS internally uses it all the time, a redirector subroutines it for FCB opens so-called "sessions" called Set FCB Owner. Clearly, this function sets the special SFI owner to the value of the current PS, but it appears to do a good deal more than that based on the behavior of MSC DEX. A redirector may have to modify the open mode before calling Set SFI Owner (see `init_sfi` in Listing 8-34 and `set_sfi_owner` in PHANTOM.C on disk).

- 0Dh Get Date And Time

This useful function returns in AX DX the current date/time in packed DDBS directory format. It can generate the times for a new directory entry generated with, for example, a subdirectory is created or when a newly created file is flushed or closed. Phantom uses this in `tbl_sfi` (see Listing 8-34).

- 11h Normalize ASCII Filename

This function translates a filename from an input buffer into an output buffer, normalizing slashes in the input buffer into backslashes in the output buffer and transferring all other characters unchanged.

- 16h Get Address Of System File Table (SFI) Entry

This function uses an SFI entry *handle* and converts it into an SFI *address*. The SFI entry number is returned in the ESI or a process at the index of the appropriate table. Since almost all redirectors store files as an SFI entry rather than a file handle, this function is quite useful with function 20h `set_sfi_owner` useful to a redirector which supports the 2Fh function (the subroutines 27h Remote Copy, which is passible file handles rather than SFI addresses).

### ■ 18h: Get Caller's Registers

This is a potentially useful function to a redirector that wants to know more about the originating DOS call. It returns a pointer to the saved register contents upon entry to the DOS function dispatcher. See the lengthy discussion of this function in Chapter 6.

### ■ 1Ah: Get File's Drive

Given a fully qualified file specification, this function separates the path from the drive letter and returns the drive number, as which it refers. It could help a redirector that uses a standard DOS drive identification string (e.g., `abc:\xyz`) as a drive letter immediately followed by a colon. The usual, for use, function could use the Phantom user's distinctive device interface that disallows use of this unsophisticated function. To determine whether the path is relative to a specific drive and not to the default drive, subfunction 1Ah works as a colon in the second character position. DOS uses it to parse paths at the INT 21h level.

### ■ 1Bh: Compare Filenames

This function compares filenames, ignoring case and treating forward and back slashes as equivalent. It could help if the redirector's underlying file system has a UNIX flavor. (Most of the rest of DOS does not.) Case is equivalent. You can even pass forward slashes to INT 21h, as do some programs such as PKZIP.

### ■ 1Ch: Make Current Directory Structure Invalid

This function invalidates the CDS associated by the DRIVE\_CDSPTR field of the SDA. To achieve this, it simply zeroes out the FILESYS field of the CDS entry. If a redirector encounters a fatal error, this function provides a simple way of invalidating its drive, since DRIVE\_CDSPTR is already appropriately set.

### ■ 20h: Get Job File Table Entry

This function, when together with subfunction 16h, provides a mechanism to convert from a file name to a job file entry. The current process to an SFT entry address, only helps a redirector that supports INT 21h function 27h. Remote copy.

### ■ 22h: Set Extended Error Info

This subfunction provides a means for modifying the DOS internal association between an extended error code and its associated class action and access fields. Phantom sets this information directly into the SDA, which appears to be more reliable than subfunction 22h when more than one redirector is present on the machine. With subfunction 22h, one redirector might override modifications made by another. MSDOS, however, uses this subfunction.

### ■ 23h: Check If Character Device

This subfunction walks the device driver header chain searching for a character device with the same specification as the FILESYS field of the SDA. Phantom uses this function to weed out and fail any calls that attempt to create a file or subdirectory with the same name as a character device such as CON or AUX. DOS does not refer to these only if no explicit path is included.

### ■ 26h: Open File

This function, together with functions 27h, 28h, and 29h, has the useful capability of performing read-only file access from within a redirector. It is interesting and perhaps disappointing that there is no Write To File function in the group. MSDOS, which of course is read only, uses these functions. This provides the means to read configuration files on the fly. This function consumes an SFT entry and a JFT entry in the current PSP.

- 27h Close File

Closes a file belonging to the current PSP. This might be a file opened using function 26h.

- 28h Move File Pointer

Performs a seek on the specified file, which belongs to the current process, and which may have been opened using function 26h.

- 29h Read From File

Performs a read from the specified file, which belongs to the current process, and which may have been opened with a submode of 26h. This call modifies the CURRENT\_Ahead of the SPA. If the call is made from within a redirector read or write call, the redirector would save the current contents of the CURRENT\_Ahead before this call and restore it afterwards.

## The Future of the DOS File System

As we said at the beginning of this section, Microsoft introduced the redirector interface in DOS 3.10. In DOS 6.0, it is still in place, little changed, and Phantom works in previous versions of DOS 7.0 from Chicago too.

However, the redirector's future looks a little confused. Microsoft, with the introduction of Windows for Workgroups, has moved from the REDIRVF-ISR line to a Windows virtual device driver, VREDIR 386, to implement their redirector. This redirector traditionally translates services to the redirector interface and the Server Message Block (SMB) protocol, on which many PC-peer-to-peer networks are based. On the other hand, a other network-related areas of Microsoft development there is apparently a strong away from the redirector interface, presumably because of the relatively poor skill set (inter- as expertise) in that area, or perhaps because of the increasing weight of the INI 21h call, in favour of an INI 21h call concept, replacement strategy that we advised against at the beginning of the section!

The gap in the network operating system marketplace, Novell has been taking the opposite approach. In NetWare 4.0, the workstation component, NETCOMMON, in previous versions, consists of several modules, headed by VMMIME, which is the Virtual-Installable Module manager. One of the modules that it manages is REDIRVIM, a file redirector. For backward compatibility, Novell provides NETVIM, which provides the traditional INI 21h replacement component functionality (see Chapter 4).

The difference between these two approaches suggests several possible alternative views. The most pessimistic is that Microsoft is in the process of phasing out the redirector interface, that VREDIR 386 simply represents a platform for increasing the support for Windows and networking, and that it will soon host an entirely different file system interface. This would imply that Novell (and the various suppliers) does not realize this and has potentially wasted a lot of time and development cost on a lame duck.

It is possible that Microsoft is indeed wary of encouraging use of the redirector interface because of its rather weak, past and backward support status, but that it will continue to tolerate it, that view would endorse Novell's strategy of making workstation file system interface modular, but it still leaves uncertainty.

One might, from a more optimistic perspective, conclude that the redirector interface is not only to remain a viable, but that through guesswork or information, Novell expects the level of support for it to increase.

Although it is part of larger changes in the implementation of the DOS file system, just as WW includes VREDIR 386, Microsoft also has VFAT 386, a 32 bit protected mode Windows virtual

dev.ac driver supporting the FAT file system. VFAT 386 is part of Microsoft's Chicago project (DOS 7.0, Windows 4.0), but it is appearing early in Win 3.11, along with VREDIR, VSHARE, and other file system VxDs. As noted earlier, Chicago will also introduce an IFSMGR as the documented, improved, and approved way of writing installable file systems.

In addition to talking a lot about the FAT directory (DPS, SEF, IF1, DPB, System FCBs, and so on), one central point emerges from this ridiculously lengthy chapter: The DOS file system isn't just for partitioned disks anymore. Any file system is primarily a logical rather than a physical construct. Interfaces such as the redirector have steadily moved the DOS file system away from the mundane world of cylinders, tracks, sectors and toward a more abstract notion of file store. A disk directory or file is anything that behaves like one.



## Memory Resident Software: Pop-ups and Multitasking

by Raymond J. Michels

With the release of MS-DOS 5.0, Microsoft documented many previously undocumented DOS functions, including ISR-related functions. Not only are these functions documented for DOS 5.0, they are also documented back to the inception of DOS in order to be just appeared. Some programmers used a "hack" but because these functions were undocumented, it was unsafe to use them. In retrospect, using these undocumented functions is not so bad because they have become safe. By using these undocumented functions, software developers pushed Microsoft into documenting and supporting them. So many of the undocumented functions are not automatic, all are safe.

There is a set of undocumented DOS calls that DOS with which PC programmers are generally familiar, but are not documented in the programming manuals. Because some programmers use the DOS Terminate and Stay Resident (ISR) functions (INT 21h Function 34h or the older ISR interrupt INT 37h), they are often called "NRs." However, these numbers are not the same as writing ISRs, but are reserved calls. INT 21h Function 34h is well known to most PC programmers, but many do not know the other DOS functions to pop up windows as part of ISRs.

Generally, programming interrupt-related memory-resident software in the DOS marketplace is not straightforward. Each has well-written DOS calls (INTs) with ISRs. Microsoft's *Using Task Link* defines the position of the software. Richard Whittom's "Termination and Stay Resident Calls" in the massive *MS-DOS Encyclopedia*, Whittom's article "Accessing File, Window, and Undocumented DOS functions and interrupts"

- INT 21h Function 34h (Return InDOS Pointer)
- INT 21h Function 50h (Set PSP Segment)
- INT 21h Function 51h (Get PSP Segment)
- INT 21h Function 534h (Set Extended Error Information)
- INT 28h (Keyboard Bust Loop)

has the following information documented in Microsoft's *MS-DOS Encyclopaedia Reference*. In addition, some vendors may provide information on the PC assembly bibliography include generic ISRs. Tables containing the complete set of DOS functions, ISRs, pop-up calls, graphics, for each issue of *PC Magazine* are also available. Online source, assembly listings and more show options show the intricacies of using these DOS functions.

Some of the functions that are the private generic ISR services are also available. These are the undocumented DOS and some extensions to the applications that include some. Anyone who is writing a SP or other code using a custom ISR library. Free and restricted binaries provide a convenient check-out mechanism through the ISR debugging environment, which could otherwise take a long time to write and test. Some of the commercial ISR libraries available are as follows:

- TeSeRact: Real-Resident Development System (TeSeRact)
- CodeRanger: for C or assembler (Microsystems Software, Framingham, MA)
- \*resident C\*: Hold Everything and TSRtic (South Mountain Software, South Orange, NJ)
- C Tools Plus 6.0: Blaise Computing, Berkeley, CA
- TSRs and More: C/C++, TurboPower Software)
- Object Professionals 1.0: TurboPower 5.5 (TurboPower Software)
- Magic TSR 1w-4w for ASM (Quantum Corporation, Cupertino, CA)
- Stay Res Plus (for BASIC, Micro Help)
- RTCCOM: batch file compiler with TSR option (Wentham Software, Wentham, MA)

In addition, and especially if Pascal is your favored language, check out the source code for Tim Kolkinger's excellent MARK and RITE ASM programs available on CompuServe (GO BPROGAS).

If there isn't any new to say on this subject, surprise!—yes. Several areas of TSR programming which our low-activated DOSs have not yet adequately covered elsewhere. These include

- INT 21h Functions 5D06h–5D0Bh (Get DOS Swappable Data Area)
- Effective TSR termination
- User C: Both Microsoft and Borland interrupt functions
- Writing non-pop-up TSRs

In addition, the Task Switching APIs provided by Windows and DOSMHLL (though undocumented) are new enough to require some mention here. We'll touch upon their use in TSRs for solving the "multiple-program-and-functio" and "double-click" switches.

Many tasks that would be TSRs in one situation should today perhaps be instead carried out by a Windows virtual device driver (VxD; see the end of Chapter 3). In addition, there are now several competing standard file loading protection tests (TSRs); these are discussed briefly in Chapter 4 (see the section on "Protected Mode DOS").

This chapter presents a variety of TSR sections that you can use to "hook" your own programs, using either Borland C++ (version 3.4x) or Microsoft C (version 6.0) as a trigger. You can use this portion of the book to integrate other parts of this book into pop-ups that are activated by the press of a user-defined function key. We discuss these items in more detail, of course. The last section of this chapter presents an interesting last program that is not activated by a key, but instead is periodically activated by the PC's timer tick, thereby maintaining in the background with whatever programs you run from the DOS command line in the foreground. This program is in addition to the PRINT multi-tasking TSR that comes with DOS.

### ***TSR: It Sounds Like a Bug, But It's a Feature***

Only three functions are absolutely necessary to write memory-resident software for MS-DOS; these three functions have been helpfully documented since their inception. They are

- Terminate and Stay Resident (INT 21h function 31h)
- Set Interrupt Vector (INT 21h Function 25h)
- Get Interrupt Vector (INT 21h Function 35h)

A TSR is any DOS program that calls INT 21h function 31h (or the obsolete but equivalent interrupt INT 27h). The description of this function's purpose in the IBM DOS 3.3 Technical Reference 5-7 terminates the current process and attempts to set the initial allocation block to the memory size in paragraphs. Does not deal with exit. The TSR function serves much like the normal DOS termination function (INT 21h Function 4Ch) which kills off whatever program calls it. The difference is that, after calling the TSR function, all memory belonging to the program is not released

Instead, part of it of the program's initial allocation block is preserved so that it will not be worked by the next program to be loaded.

Thus, a TSR's save DOS program that saves bits of itself behind after terminating (this sounds like a class, but sometimes referred to as the "leaky bucket") where memory is allocated and then is cyclically located. It doesn't sound like a feature you would want in your operating system, but one area in which in-house software subindustry could be built.

What is the advantage of terminating without freeing your memory. If you've terminated a real some other program's memory, there's not much your memory is going to do other than take up space until it's being up memory can occasionally serve a purpose. I know. TSRs have been written with names like Mr. MHO and I MEM to allow a developer with say a 640K machine to test software under conditions similar to those on say a 512K machine. Because from this limited use, what good is it to log memory after you're gone. You can't take it with you!

How, so what, the second necessary function, Set Interrupt Vector, comes in. All functions based on the Intel 80x86 architecture allow any program to create code that gets invoked whenever a hardware or software interrupt is generated. For example, the only reason IN+ 21h is a gateway to MS-DOS services is that the interrupt vector 21h points to a code inside DOS that provides the services. The ability to hang a piece of code off of an interrupt vector is what makes the TSR function something other than a real obscure way to conserve memory. You can use the Set Interrupt Vector function to point interrupt vectors to your code, and then call the ISR function to keep this code and its associated data in stack space resident in memory after you terminate. Whenever one of your interrupts is generated, it executes the code you left behind. Thus, theoretically, it's left after termination, you can take it with you.

What sort of interrupts would a TSR be interested in trapping. The most obvious one is the hardware interrupt INT 9, generated every time a user presses a key. By trapping INT 9, a TSR can watch every key that a user types. Let's say our TSR is memory resident under and N-solar computer that passes windows from *The Mabinogion* when the user presses Alt-M, to *Paintout Paintout* when the user presses Alt-P. These are the only two keys the ISR is interested in, and they are referred to as the program's hotkeys. Each time the user hits a key, the INT 9 hardware interrupt looks at a key, and if it is not one of its hotkeys, goes back to sleep. But if it is one of its hotkeys, then your application should do its thing. In this example, that means paying eight opeas "Largaria," but in a TSR it would this sudden we're springing a trap, it's called, by popping.

Now, one thing has been glossed over. When the user types a key that is not one of the ISR's hotkeys, how does the key go to its destination. The ISR can't disregard it, so it must somehow let other programs get a crack at it. The ISR does this by jumping to whichever memory previously saved the INT 9 vector before our ISR installed its INT 9 handler. Another way to look at setting an interrupt vector is most all, TSRs have to get the interrupt vector's previous value, calling the DOS key Interrupt Vector function. Thus, the ISR looks something like this:

```

INTERRUPT PTR old_int9_handler;
INTERRUPT my_int9_handler()
    IF (key == alt_m)
        m_sadof();
    ELSE IF (key == alt_p)
        p_sadof();
    ELSE
        JMP PTR old_int9_handler();
BEGIN
    old_int9_handler = _dos_getvect(9), // INT 21h Function 35h
    _dos_setvect(9, my_int9_handler), // INT 21h Function 25h
    go_tsr(); // INT 21h Function 31h

```

Every program that has hooked INT 9 takes care to call the interrupt's previous owner, then every program that needs to will get a peek at the stream of user keystrokes. Jumping to the previous owner's job will also end the interrupt, and the end result is an interrupt chain (see the INTERRUPT CHAIN program at chapter 6). Every time you press a key, a whole host of programs might see it. This mechanism for multiprogram access to the keyboard input stream was formalized in the OS-2 concept of the "monitor."

There is a structural difference between CALLing the previous interrupt handler and JMPing to it. Handlers that want to process an interrupt *after* the previous handler (i.e., postprocessing) pass control to the previous handler with a CALL. Otherwise, the handler pre-processes the interrupt and passes control with a JMP. Unlike a CALL, a JMP of course, does not return.

While it's best known ISRs, such as Borland's now ancient Subkey, are pop-ups that are activated by hitting a key, pressing a hotkey is just one way of generating an interrupt. Anything that generates a interrupt can activate an ISR. For example, when our own program calls INT 21h it's generating a software interrupt so an ISR could easily attach itself to INT 21h, providing a mechanism for stealing the operating system or for debugging, as in the INTERSPY ISR in Chapter 8. For example,

```

INTERRUPT PTR old_int21_handler,
INTERRUPT my_int21_handler()
IF ($ah some function I'm interested in)
    // do pre processing
    // maybe CALL PTR old_int21_handler()
    // do post processing
ELSE
    JMP PTR old_int21_handler(),
BEGIN
old_int21_handler = _dos_getvect(Dx21);
dos_setvect(Dx21, my_int21_handler),
go_isr ).

```

In this example, as soon as we attach my handler to INT 21h by calling \_dos\_setvect, the processor passes all INT 21h calls to the code in my handler. This means that our own call to INT 21h (function 31h) will pre-process itself first processed by my handler. It will then jump to the code in my handler to do my work, and afterwards call INT 21h request. Presumably, the call to function 31h would pass by another computer to old\_int21\_handler, which might be MS-DOS or some other ISR that pre-processes INT 21h stuff, as in the simple application wrapper's from the end of chapter 2.

With all this pre-processing, it's not clear that programs receive the ISR benefits for genuinely useful code that's worth having, code that's necessary. Software that helps the user prepare to go to bed and tests to see if a computer is not a good candidate for memory residency. Neither for that matter is our Courier and Sun Microsystems since it's fully available, dedicated hardware already exists for this purpose.

## Where Does Undocumented DOS Come In?

Not a many functions you need to produce the ISR are fully documented, where does undocumented DOS come in. Does it really need undocumented DOS to write a program that plays "I Am the Way Model of a Mx5x" Minor General, when all the user presses the A/P hotkey. Unfortunately, you are not able to do it unless you have achieved remarkable data compression, you don't want to, not so for the memory occupying memory. Instead when the user presses A/P you want to allow some memory, and the music in from a file, cause the file plays the notes, free the memory, then go back to sleep.

It would be nice if things worked this way, but they don't. The problem is that in this example you have no control over when my\_int9 handler will be invoked. Recall that my\_int9 handler is not called from within the program, the way functions like \_dos\_setvect() or go\_isr() are. Instead

my 109 number) is called whenever the user pounds on the keyboard. A keypress is an asynchronous event that bears no relation either to the internal state of whatever program happens to be running or to the internal state of DOS. You can't control when the user will press a key.

For instance, the foreground program might be copying a large file to the printer when the user decides there's time for a musical interlude. If while the foreground program is executing an INI 21h function such as Read File or Write File, the penultimate function suddenly takes over and starts issuing its own INI 21h requests, the resulting scenario is one which DOS was not designed to handle. MS-DOS is a single-tasking operating system, which means that it does not allow for the possibility that it might be interrupted in the middle of one request, be asked to carry out some other request, and then resume the first request at the point where it was interrupted.

This property of MS-DOS is often referred to as "non-recentrancy," meaning that if INI 21h is already executing, another INI 21h request can't be issued. A function or program which is "reentrant" is capable of so that it can be interrupted at any time, allowing another process to enter without losing the state of the function just before the interruption. There are many techniques to achieve reentrancy; a function will normally keep its state in the contents of variables that exist to a particular instance of its invocation on the caller's stack. With a very few exceptions, MS-DOS uses neither this nor any other reentrancy tricks. As we saw in chapter 6, DOS relies almost entirely on global data tables, which uses its own stacks, rather than the caller's.

Any textbook on operating systems or on concurrent programming contains a discussion of the difference between reentrant code, which may be started by several processes simultaneously, and what is known as called serially reusable code, which may be used by only one process at a time. MS-DOS is serially reusable code.

When MS-DOS is called using INI 21h, DOS searches as one of three internal stacks: the I/O stack (or Disk stack or the Auxiliary stack), functions 00 through 0F, or the I/O stack. The remainder of the time consists of the Disk stack. If MS-DOS is called for a critical routine, such as DIR A, when the drive door is open, the Auxiliary stack is used. Incidentally, we saw exactly how DOS uses these stacks. Because of this stack switching, if a TSR calls MS-DOS when the foreground is already executing inside INI 21h, MS-DOS loads the TSR's data onto its stack, overwriting the foreground process's data.

If DOS is impossible servicing a function 0Fh request or lower-numbered I/O requests, a function 0Fh request or higher than there would be a problem because two different stacks are involved. Furthermore, as we again know from chapter 6, a few INI 21h functions—33h, 50h, 51h, 62h, and 64h—use stacks that they use the caller's stack and as therefore fully reentrant. But for the most part, DOS is non-reentrant.

It's not completely true, since it's not for our DOS you have to worry about. What if the heads on the hard disk are in the middle of writing data as part of the response to an application's INI 13h call. If the disk starts issuing INI 13h requests that move the heads somewhere else, then you're going to have a serious reentrancy problem that has nothing to do with stacks or reusable code, but that could well result in a scrambled hard disk.

Does this mean the TSR can't perform DOS memory and file operations whenever the user presses the keys? Does it thus have to be done once during initialization, before hooking any interrupt vectors, so that the memory resident portion of the TSR avoids all use of DOS calls? For example, one source of programming for the PC makes the blanket statement that a TSR that requests service outside I/O "cannot use any DOS functions." If this were true, I would certainly restrict what you can do with TSRs.

This sort of question is to be a restriction as it sounds. Many commercial programs for the PC that aren't in the I/O bypass DOS for many operations such as screen display and keyboard input. Avoiding DOS is so far as possible, so for certain key operations on the PC, it is practical to bypass it. It is easy to write screen display functions, for example, that not only bypass DOS I/O, which are again

times faster than DOS output routines and which provide far greater control over the screen. Not being able to receive DOS sounds almost like a blessing in disguise.

However, as we noted in Chapter 8, the one area of DOS's functionality that nearly is irreplaceable is its I/O. For a lot of "wild" programs can allocate expanded or extended memory rather than use the DOS memory allocation routines; expanded or extended memory is not always available, so many ISRs use local static memory through DOS. In summary, most ISRs need to make some INT 21h calls while popped up.

Fortunately, it is simply not true that ISR interrupt service routines "can't make INT 21h calls." But it is true that ISRs must do something special to make such calls. There are two options:

- Make some INT 21h calls while INT 21h is in the middle of processing a request, or
- Somehow save and restore all of DOS's context, including the three DOS stacks, so that you can freely interrupt it.

The second option will be discussed later in this chapter, in the section on the DOS Swappable Data Area (SDA). And then we will concentrate on how not to enter DOS in the middle of some other activity, i.e., an INT 21h call, but instead to wait until that call has completed—how to use DOS as a service task, i.e., a service. Until we discuss the SDA, you will be reading about the state that a ISR must save and restore as part of its pop-up regime.

The critical requirement here is to have some way of determining when INT 21h is busy or, more accurately, of determining when any of its three stacks is in use. A short while ago, you saw a small program (see Example 1) for trapping INT 21h calls, and it may have occurred to you that this might help determine whether DOS is in use. For example, you might put both the INT 21h handler and INT 9 handler into the same program and, so the former to tell the latter whether it's safe to pop up.

```

INTERRUPT PTR old_int9_handler, // keyboard
INTERRUPT PTR old_int21_handler, // DOS
WORD using_io_stack = 0;
WORD using_disk_stack = 0;

INTERRUPT my_int21_handler()
  IF (ah <= 0x0c)
    INCR using_io_stack;
    CALL PTR old_int21_handler;
    DECR using_io_stack;
  ELSE
    INCR using_disk_stack;
    CALL PTR old_int21_handler;
    DECR using_disk_stack;

INTERRUPT my_int9_handler()
  IF key == alt_m AND NOT using_disk_stack
    mikado();
  ELSE IF key == alt_p AND NOT using_disk_stack
    penance();
  ELSE
    JMP PTR old_int9_handler();

BEGIN
  old_int9_handler = _dos_getvect(9);
  old_int21_handler = _dos_getvect(0x21);
  _dos_setvect(0x21, my_int21_handler);
  _dos_setvect(9, my_int9_handler);
  go_isr();

```

Each hook calls INT 21h to find out whether someone is "in DOS." The INT 21h handler increments a flag on entry to an INT 21h call, and decrements it on the way back out. The INT 9 handler checks the using\_disk\_stack flag and won't pop up if the flag's non-zero. The flag therefore acts as a semaphore, serializing access to DOS.

This is the basic idea behind making DOS calls from a TSR. But there are many problems with the preceding pseudocode. For example, DOS termination functions (functions 00b, 31h, and 004h) do not return, and therefore would need to get special treatment. Likewise, raw code does not take into account DOS critical errors. Nor does it account for the fact that PC strings (i.e., COM, MAND.COM) prompts is actually parked inside INT 21h function 04h. Bothered as I am, I spent making it seem as if one can't pop up while in the DOS prompt, which you already know is not the case. Fortunately, this code doesn't need to work properly because MS-DOS already provides with an InDOS semaphore and a critical error semaphore that the TSR can check. Instead of brooding INT 21h in an attempt to maintain the InDOS flag, you can use the one that DOS already provides. On the other hand, this technique of brooding an attempt to maintain an InDOS semaphore is out to serialize access to INT 13h (i.e., ROM BIOS) disk interrupt.

This is where newly documented DOS enters the picture, because the DOS function that returns the address of the InDOS semaphore was once undocumented, and the documentation still doesn't reveal one how to use this function—also called by a critical error semaphore. Furthermore, DOS generates the interrupts INT 28h while inside INT 21h function 04h. As you can see later, Microsoft originally created these workarounds for its own TSRs, such as PRINT.COM.

Like Microsoft's own TSRs use these functions to avoid a TSR that developers who wear a corporate to just TSRs probably need to use them as well. It goes against common sense to assert that using undocumented features makes a program more rather than less stable, but who said that TSR programming was supposed to make sense. The techniques for writing correct TSRs may not be a model of software engineering at its finest, and some of the undocumented functions for TSR support have been used to generate alert messages earlier than parts of a well thought-out interface, but you need them if you want your program to survive in the PC marketplace.

If you're writing a TSR, you have probably already brought into a host of computer problems, and I'm only using undocumented DOS as the least of them. The reason is fairly plain: I have spoken of a "TSR error" not because of undocumented DOS, but because of keyboard conflicts, problems associated with popping up over screens or graphics mode, memory usage conflicts, and the like. Undocumented DOS is one of the same areas of TSR programming. This is part is confirmed by the fact that Microsoft finally documented a number of the TSR necessary functions that were previously undocumented.

## MS-DOS TSRS

How did PC programmers find out about the undocumented TSR functions? From examining Microsoft's own TSRs, of course.

TSRs have been a part of MS-DOS since its first release in 1981. Microsoft documents in its first article, "Sub Memory Resident Programming," *The Wall to Wall MS-DOS Digest*, 1988, that TSRs were first available within the 64K confines of the original PC-M operating system in programs like SmartKey, Uniform, and Unsplit. The only TSR programs to ship with DOS 1.x was MD02.COM, PRINT.GRAPHICS, and ASCII.S were first in DOS 2.x.

PRINT is one of DOS first program that preempts external tasks. You can run an application at the same time that PRINT is printing a file. Only one of the two programs is running at any given instant, but the illusion of simultaneous operation is maintained by switching between them on a timer tick.

When the PRINT program is installed, it chains into the DOS timer tick interrupt (INT 1Ch) by the DOS Keyboard Bios Loop interrupt (INT 28h) and generates the interrupt by taking a large number of interrupts as the normal for a TSR. INT 1Ch and INT 28h allow the PRINT program to get control at regular intervals independent of the user and to perform its processing operations in a predictable cause. These intervals are sufficiently close together that you might not program appears to be operating at the same time as the PRINT program.

Every time PRN1 wakes up, it saves the current DTA, PSP, and the vectors for INT 13h (Ctrl Break), 23h (Ctrl C), and 24h (Critical Error). PRN1 sets up its own values for these items. On exiting, it restores the current processing. PRN1 restores these values to their original state. The TSRs presented in this chapter follow a similar structure. The introductory TSR at the end of this chapter is an enhancement to the PRN1 TSR. It periodically looks for files that appear in a certain subdirectory and automatically submits them to PRN1.

One of the significant improvements to MS-DOS 2.0 was the availability of a hard disk. However, this new disk-only drive (C:) did cause some problems with software that was hard-coded to use drive A as the root. The ANSI.SYS utility allowed drive letters to be mapped to other drive letters. When programs refer to root drive A, they could be transparently made to access drive C, instead. ANSI.SYS version three (MS-DOS 3.0) supports INT 21h (DOS Function Call), INT 25h (Absolute Sector Reads), and INT 26h (Absolute Sector Writes). When INT 25h or 26h is called and the M register references the ANSI.SYS device, the value of M is replaced by the new drive number. Sample lists.

MODE is another example of a TSR that causes output to a device. Many programs do not support external printers (COM1) or just reference EPP1. Among other capabilities, MODE can generate text through INT 17h (BIOS Parallel Printer Service) and send the data to the serial port. The same principle can be used to write translators for programs for various output devices. A TSR program is one of the authors' once, wrote way for a printer that did not support the form-feed command. The DOS could be passed printer output interrupt and checked for a form-feed character. When one came, the TSR would setup the appropriate number of line feeds to get to the next page. It was a simple program, but it saved someone from having to buy a new printer.

## The Generic TSR

Now it's time for some code! Our goal is to build a generic TSR with Microsoft or Borland C/C++ compilers. This chapter provides everything you need when a user-defined hook is pressed, a function named `__pfnhook` is called. You simply provide application-specific functions with the generic TSR object modules, and you're done. We've used several strategies to avoid some annoying screen modes or even screen save and restore, since these have not always done well in undocumented DOS. The generic TSR code deals with all the tasks connected with undocumented or once-undocumented DOS. The resulting TSRs have been tested extensively and seem quite robust. There are no 100% guarantees in the world of TSRs, however.

We'll use the generic TSR to build three sample programs: a simple file browser (TSRFILE), a memory resident version of the MBROWSE from Chapter 7 (TSRMBROWSE), and a memory resident version of the INT 21h command interpreter that we will meet in Chapter 10 (TSR21). We also build MAKEFILE, the one-step autoinstalling program mentioned earlier. The TSRs can be built either using the traditional DOS technique for TSRs or by using the DOSXWAP technique described later on. Finally, we will discuss whether a given TSR uses the disk or not. To show some of these pieces together, we'll take the source code as our approach, it best showing the MAKEFILE for this project. Source Listing 9-1 (in the book) works with NMAKE from Microsoft C 6.0. Make files for Borland C++ are also supplied on the diskette.

### Listing 9-1 MAKEFILE for the generic TSR

```
# MAKE makefile for generic TSR
# example C:\UNDOC2\CHAP9>make tsrfile.exe
# can be overridden from environment with MAKE /E
# example,
# C:\UNDOC2\CHAP9>set swap=1
# C:\UNDOC2\CHAP9>make /E tsrfile.exe
#
```



```

# C:\UNDOC2\CHAP9>set no_disk=1
# C:\UNDOC2\CHAP9>make /e tsrmem.exe
SWAP = 0
NO_DISK = 0
!IF $(SWAP)
DOSSWAP = -DOS_SWAP
!ELSE
DOSSWAP =
!ENDIF
!IF $(NO_DISK)
USES_DISK =
!ELSE
USES_DISK = -DOSES_DISK
!ENDIF
# defines the key components of the generic TSR
# TSREXAMP.C - main
# INDOS.C - IndOS, critical error flag
# PSP.C - Set PSP, Get PSP
# EXTERR.C - Extended error save and restore
# TSRUTIL.ASM - Miscellaneous routines
# STACK.ASM - Stack save and restore
# DOSSWAP.C - Optional use of DOS Swappable Data
# SWITCHER.C - Task switcher, instance data
# NOTIFY.ASM - Task switcher notification handlers
# BREAK.C - Ctrl-C handling
UNDOC_OBJ.S = indos.obj psp.obj exterr.obj break.obj switcher.obj
MULTI_OBJ.S = indos.obj psp.obj exterr.obj break.obj
TSR_OBJ.S = tsrexamp.obj $(UNDOC_OBJ.S) \
    tsrutil.obj stack.obj notify.obj
$TSR_OBJ.S = tsrexamp.obj $(UNDOC_OBJ.S) dosswap.obj \
    tsrutil.obj stack.obj notify.obj
# command to turn a .C file into an .OBJ file
.c obj:
    cl -AS -Ox -Zp -c -W3 -Zi -DTSR $(USES_DISK) $(DOSSWAP) $*.c
# command to turn an .ASM file into an .OBJ file
.asm obj:
    masm -ml $* .asm,
# special handling for MULTUTIL.ASM
multutil.obj: tsrutil.asm
    masm -ml -DMULTI tsrutil,multutil;
multitask.obj: stack.asm
    masm -ml -DMULTI stack,multitask;
# make the file-browser sample TSR
tarfile.exe: $(TSR_OBJ.S) file.obj
    link /far/mol/mop $tarfile.rsp
# make the RCB-walker sample TSR
tsrmem.exe: $(TSR_OBJ.S) mem.obj put.obj
    link /far/mol $tsrmem.rsp
# make the INT 2Eh command interpreter sample TSR
INT2E_OBJ.S test2e.obj send2e.obj have2e.obj do2e.obj
INT2E = test2e send2e have2e do2e
tsr2e.exe $(TSR_OBJ.S) $(INT2E_OBJ.S) put.obj
    link /far/mol $tsr2e.rsp
# link the dos swapple examples only if the env var set
!IF $(SWAP),
# make the file-browser sample TSR
starfile.exe: $(SYSR_OBJ.S) file.obj

```

```

link /far/nop @starfile.rsp
# make the MCB-walker sample TSR
starwem.exe $ (STAR_OBJ) mem2.obj put.obj
link /far/nop @starwem.rsp
# make the INT 2Eh command-interpretor sample TSR
star2e.exe $ (STAR_OBJ) $(INT2E_OBJ) put.obj
link /far/nop @star2e.rsp
ENDIF
# make the non-pop-up PRINT add-on
multi.exe mu ts obj $(MULTI_OBJ) multitut1.obj multstk obj put.obj
link /far/nop/mep/1 multi $(MULTI_OBJ) multitut1 multstk,multi

```

## TSR Programming in Microsoft and Borland C/C++

When beginning our examination of the various components of the generic TSR, we need to discuss a little TSRs in Microsoft and Borland C++ rather than in assembly language. This discussion strays a bit from the scope of understanding DOS architecture, but that's unavoidable. As a consequence, we've covered a sort of interesting aspects of low-level PC programming in C.

Managing SRVs in assembly language seems relatively easy at first because you have total control of the CPU. The process becomes a bit more difficult when the actual TSR application goes beyond the scope of simple assembly language code. C is generally easier to code than assembly language, and much more maintainable. By using C to write a TSR, you give up a little efficiency but gain ease of use and manageability.

For a TSR to work on time it must be accessed through some type of interrupt. Therefore, anyone interested in TSR programming in a high-level language like C must become familiar with the facilities for interrupt manipulation.

Most C compilers for the PC come with an interrupt or interrupt keyword that helps create interrupt handlers and ISR's. The interrupt keyword causes the compiler to create special entry and exit code on the processor whose behavior is like interrupt modules. On critical functions save each of the registers and save DS to that of the C program. Because these registers are defined as parameters and a common set on the stack, you can get and set them just like any other variable. When the procedure exits, it pops the registers values from the stack. Listing 9-2 shows example code for a simple interrupt handler.

### Listing 9-2: An interrupt handler in C

```

typedef struct {
#ifdef TURBOC
    unsigned bp, di, si, ds, es, dx, cx, bx, ax;
#else
    unsigned es, ds;
    unsigned di, si, bp, sp, bx, dx, cx, ax; /* PUSH4 */
#endif
    unsigned ip, cs, flags;
} INTERRUPT_REGS;

void interrupt far my_handler(INTERRUPT_REGS r)
{
    unsigned i = r.ax;
    r.bx = i >> 8;
}

```

Sample code for the interrupt keyword shown shows an enormous parameter list for each interrupt handler with each register named separately. Using the INTERRUPT\_REGS structure, not a pointer to one, makes the parameter list more manageable.

What sort of code does this produce? By compiling with the Microsoft C 5.0 or 6.0 command line switches, you can examine the resulting assembly language code. Listing 9-3 shows the code generated

by the compilation of this interrupt procedure in Microsoft C. The order in which registers are pushed is dictated in part by the Intel PUSH and POPA instructions. For and C++ also offers an interrupt keyword, but the order in which it pushes and pops registers is different from and incompatible with the PUSH/POPA instructions.

### Listing 9-3: Assembly Language Generated from Listing 9-2 by Microsoft C

```

_my_handler PROC FAR
    push    ax                ; bp+18
    push    cx                ; bp+16
    push    dx                ; bp+14
    push    bx                ; bp+12
    push    sp                ; bp+10
    push    bp                ; bp+8
    push    si                ; bp+6
    push    di                ; bp+4
    push    ds                ; bp+2
    push    es                ; bp+0
    mov     bp,sp
    sub     sp,2
    mov     ax,DGROUP
    mov     ds,ax
    ASSUME DS DGROUP
    cld
    mov     ax,WORD PTR [bp+18] ; i = r ax
    mov     cx,ah
    sub     ah,ah
    mov     WORD PTR [bp+12],ax ; r bx = i >> 8
    mov     sp,bp
    pop     es
    pop     ds
    pop     di
    pop     si
    pop     bp
    pop     bx
    pop     bx
    pop     dx
    pop     cx
    pop     ax
    ret
_my_handler PROC FAR

```

Pushing the registers on the stack allows the C function to access the register values through variables. In a sense, these values are popped from the stack so even the C function can change the actual values of registers via interrupt call. Notice that `ds` and `es` is popped twice. Once, `es` was pushed at this point. If the C function was allowed to change `SP` (the stack pointer), the IREI instruction would push `es` twice, once `es` is popped, and then IREI uses the stack to return to the caller. Notice that the processor itself pushes `CS` and the flags on the stack.

It is a compile for 80286 and higher machines with the `-G2` switch, the resulting code would look like Listing 9-4.

### Listing 9-4: Assembly Language Generated From Listing 9-2 with `-G2` switch

```

.286
_my_handler PROC FAR
    pusha                ; push ax,cx,dx,bx,old_sp,bp,si,di
    push    ds
    push    es
    mov     bp,sp
    sub     sp,2
    mov     ax,DGROUP
    mov     ds,ax

```

```

ASSUME DS:GROUP
cid
mov ax,WORD PTR [bp+18]
mov al,ah
sub ah,ah
mov WORD PTR [bp+12],ax
mov sp,bp
pop es
pop ds
popa
; pop di,si,bp; skip sp, pop bx,dx,cx,ax
ret
my_handler ENDP

```

The enormous amount of code generated (even for the best case, with PL MIA, PCPA, and the lack of use of stack space used for our three-line interrupt handler) should not go unnoticed. If you were writing this assembler in assembly language, it might look like Listing 9-5.

### Listing 9-5: Interrupt Handler in Hand-Crafted Assembly Language

```

_my_handler PROC FAR
    mov bx, ax
    shr bx, 8
    ret
_my_handler ENDP

```

Every feature has a price, and here that you pay for the convenience of writing the application in C rather than assembly language. Remember that for each interrupt your TSR hooks, every application which generates that interrupt will send you your TSR's interrupt handler, even if the application isn't registered in the TSR's services, and even if the TSR is just going to chain to the previous handler. As we saw in chapter 6, this is an especially bad problem with the long INT 21h chain. Adding a global vector interrupt handler to this chain just makes things worse.

To add to the misery, the interrupt or interrupt keyword C compilers for the PC generally offer a set of functions for manipulating interrupts. In Microsoft and Borland C, C++ the DOS.H header file provides a large set of DOS-specific functions, including those excerpted in Listing 9-6.

### Listing 9-6: Excerpts from Microsoft C DOS.H

```

void _cdecl _interrupt_far *
_cdecl _dos_getvect(unsigned intno)();
void _cdecl _dos_setvect(unsigned intno,
void _cdecl _interrupt_far *new_handler)();
void _cdecl _chasm_intr(void (_cdecl _interrupt_far *target)());
void _cdecl _dos_keep(unsigned retcode, unsigned newsz);

```

The custom \_dos\_getvect and \_dos\_setvect directly translate into calls to INT 21h functions 25h and 25h, and as you're pickable to using the more general intdosx or int86x functions, for example:

```

#include <dos.h>
//
extern void interrupt far my_int21_handler(); // declare new function
void _cdecl _interrupt_far *old_int21(), // pointer to old function
main()
{
    old_int21 = _dos_getvect(0x21); // save old
    _dos_setvect(0x21, my_int21_handler); // install new
    // -
    _dos_getvect(0x21, old_int21); // restore old
}

```

You can do more with the old `int21` function pointer than just restore it when you're finished. In fact, almost all interrupt handlers and ISRs need to do something *else* with the pointer to the previous handler. They need to chain to it! Microsoft and Borland provide the extremely useful `chain_int` function, which is necessary when your new interrupt handler needs to do preprocessing before chaining to the old handler. For instance,

```
void C_interrupt_far *old_int21();
void interrupt_far my_int21_handler(INTERRUPT_REGS r)
{
    // do some work
    chain_int(old_int21),
    // never reached!
}
main()
{
    old_int21 = dos_getvect(0x21),      // save old
    dos_setvect(0x21, my_int21_handler), // install new
}
```

It's called *interrupt chaining* because each interrupt handler is like a link in a chain, each one is a unit, but the units are linked together. The `chain_int` function is basically a `JMP` instruction. Control is passed from the current interrupt handler to the one passed as a parameter to `chain_int`. When the final interrupt handler is reached, there could be many handlers linked together. The handler performs an interrupt return, passing control back to the original application that was active when the interrupt occurred.

If you need to do more work after chaining to the old handler, called *post-processing*, then you can use `chain_ptr`. Instead you must directly call through the saved function pointer:

```
// maybe do some preprocessing
(*old_int21)();
// we're back - do postprocessing
```

The C compiler turns `chain_ptr` call through the macro `func_ptr` into the following:

```
pushf
call dword ptr old_int21
```

The problem with this, however, is that the compiler uses the CPU registers a way that was not obvious from an examination of your C code. The registers on entry to the old interrupt handler may therefore not be correct. This is not a problem with `chain_ptr` because it is not *you* which can only be correctly called from within a interrupt function, only by the CPU registers with the image of the registers that were stored in the stack. Whenever possible, use `chain_ptr` rather than `*old`.

There are many tradeoffs involved in writing interrupt handlers in C rather than in assembly language. All of them seem like a win, but the overhead of pushing old registers on the stack on entry to a interrupt handler, and the inconvenience of not knowing the exact state of the registers before chaining to the previous handler, are sometimes too much. For instance, in C it is less convenient to save a way switch state in assembly language. The generic ISR uses two separate language modules: ISRU (THE ASM) and ISMC (ASM). Because that made more sense if it was arbitrary C only principles.

### Keeping a C Program Resident

The hardest thing to do from a ISR in C is estimate the amount of memory you want to allocate. Both Microsoft and Borland provide a handy `dos_keep` function declared in `DDOS.H` (see the excerpts in Listing 9-6) which calls the DDOS TSR helper. But this does not answer all the

system of what number to pass it as `memory.c` does keep. When coding in assembly language, you can work up with this number fairly easily because you can find the size of your code. You have to be careful, however, at a higher level, allowing you to perform such efficiencies as placing startup code at the end and returning it when going remote. You can divide addresses at the end of each segment to find the address to calculate the size of the segment. You can then total the size used in each segment to obtain the size required.

Figure 9-7 shows how you do not control the memory structure of the program beyond your source code. The memory map for a small model Microsoft C program with hypothetical segment addresses is shown in Figure 9-7.

**Figure 9-7: Memory Map of a Small Model Microsoft C Program**

1E20h	FAR HEAP
	-----
	NEAR HEAP
	STACK
0E19h	DS
	-----
0BFah	CODE
	-----
0BEAh	PSP
	-----

For Borland C++, the STACK and NEAR HEAP are reversed. Several readers did not agree with this memory map, but it can be validated by running a C program in a debugger. If a break point is set at the `main` function, you can look at the value of the SP register. For Microsoft C, SP is a low address. For Borland C++, SP is at the top of the stack segment.

One note you could not pass every high number to `dos keep`, but with EMRS, one of the primary goals is to keep memory as close to optimal as the absence of a source. The stack fragment in Listing 9-8 details one way to keep a few memory segments resident in DOS or real memory models.

**Listing 9-8: Sample Code for Keeping Memory Segment Resident**

```
#define PSP_ENV_ADDR    0x2c /* environment address from PSP */
#define STACK_SIZE     8192 /* must be 16 byte boundary */

#define PARAGRAPHS(x)  ((FP_OFF(x) + 15) >> 4)

char far *stack_ptr; /* pointer to TSR stack */
unsigned memtop, /* number of paragraphs to keep */
//
/* MALLOC a stack for our TSR section */
stack_ptr = malloc(STACK_SIZE),
stack_ptr += STACK_SIZE,
//
/* release environment back to MS-DOS */
FP_SEG(fp) = _psp,
FP_OFF(fp) = PSP_ENV_ADDR,
_dos_freemem(*fp),
/* release unused heap to MS-DOS */
/* All mallocs for TSR section must be done in TSR text */
segread(&aregs),
memtop = sregs.ds + PARAGRAPHS(stack_ptr) - _psp,
_dos_setblock(memtop, _psp, &dummys),
_dos_keep(0, memtop);
```

First, create a block of memory in the near heap as `stack`. This becomes the TSR's stack during activation. The alternative to using this `stack` is to use whatever stack happens to be in

effect during TSR activation. This is fine for small programs. But if you are doing wild and wonderful things with your code, it is best to create your own stack to avoid overflowing the foreground's stack. The stack size is added to the stack pointer variable to get to the bottom of the stack. This stack bottom becomes the top of the TSR in memory.

Use the value of the new stack pointer to calculate the number of paragraphs that must be kept resident in memory. The expression `PARAGRAPHS(stack_ptr - 0x0000)` is the number of paragraphs in the local heap. This number must be retained because it includes the address you've already done. This is added to `DN` to establish the top of the memory to be freed, and the `SP` is subtracted to find the actual number of paragraphs needed by the `console` program. Call `MSDOS` to shrink the current blocks down to the size specified. In simple programs created with the `gpc` `TSR`, memory is generally less than 6000 paragraphs, giving the resulting `TSR` an in-memory footprint of about 24k. This eliminates any far heap, the original stack, and a console heap.

Using this method, you must perform any near calls before creating the stack. Once the `TSR` is resident, I cannot call the `malloc` family of `crx` library routines from the `malloc` functions because the near heap is gone. Use of `malloc` calls to library functions varies with the compiler/operating system, so be sure to select your functions carefully. The *Run-Time Library Reference* for Microsoft C 6.0 includes in the entry for `malloc` a list of all linker options that call `malloc` if you're large. The final step is to call `_dos_keep`, which does an `IN+` 2H function `51` to terminate and set resident, retaining in memory the number of paragraphs specified. If all goes well, you should be able to find your `TSR` in the display from Chapter 7's `UDMM` program. For example,

```
C:\JRD0C2\CHAP9>tarfile -k 59 4
Activation: CTRL SCAN=59
C:\JRD0C2\CHAP9>udmem
Seg      Owner    Size
---
0BEA    1E76    0000 ( 208)
0BE8    0000    0000 (  0)    free
0BE9    0BEA    0597 ( 22896)    -k 59 4 [08 09 13 28 2F ]
```

The `MEM` display shows that the `TSR` begins at `0BEA` (which is `0BE8`, of course, at `0BE9`) that it retains 0597h paragraphs (22k) and that we freed the current block. You can see the `current` field, which by the way, designates a forked off `CUI`. But the program name is missing, available. As for the various interrupts, `UDMM` says you've hooked `05`, so will be discussed later on while.

One final note about staying resident `tsr` C programs: to reduce their memory footprint even further, many `TSR`s perform their startup code (for example, `msdos`) need not run once you're gone resident. Any subsequent calls to `TSR` will be considered, for instance, as going to go to the main and a copy of a different instance of the program, not the main of the resident copy.

This is one program, but possibly more than one process. Anonymous it would be nice to throw away the code's main. This is a discussion that is relatively easy using assembly language, figuring out how to do it in C, given the memory map shows earlier, is left as an exercise for the reader. Don't stay up too late!

## Not Going Resident

If you are writing your own `TSR` from scratch, rather than using one of the commercial `TSR` libraries, or a generic `TSR` such as the one we present here, it is a good idea to put off going resident in as long as possible. Don't try disguising your application as a `TSR`. Instead, have it spawn a command shell from which you can exit, or have it run a single program whose name and arguments appear on the `DOS` command line, as shown in Listing 9-9.

**Listing 9 9: Sample Code for Shell vs. TSR Implementation**

```

main(int argc, char *argv[])
{
    // TSR init goes here
    old_int09 = _dos_getvect(0x09),
    _dos_setvect(0x09, my_int09_handler),

#ifdef TESTING
    // to launch a command shell
    system(getenv("COMSPEC")),
    // or, to run just one program:
    // spawnvp(P_WAIT, argv[1], &argv[1]),
    // we're back, deinstall
    _dos_setvect(0x09, old_int09);
#endif
    //
    //
    _dos_keep(0, memtop),
}
}

```

In fact, it is now fairly straightforward to consider making some of your applications into shells rather than TSRs. A program that needs to set up a context of some sort for another program is often best treated as a shell or a TSR. See the two application wrapper examples at the end of chapter 2 (DOSVER and FLSM0132).

**Jiggling the Stack**

Recall that the stack we created with `make` just before remaining resident. For that stack to be used, the TSR interrupt routine that performs activation must call two routines: one sets up for local stack on entry and one restores the original stack on exit. The actual code to switch stacks must be programmed in assembly since we need to have full access to the registers (it is though you could use the `int` assembly instruction to do indirect register redirection—function of `PHANLOM.C` in Chapter 8).

Listing 9 10 shows the assembly language module that manages the stack context switch. The `_set_stack` procedure saves the current foreground stack in the `ds` area and sets the stack pointer to the stack created with `make` by `stack_ptr`. Notice the stack manipulation at entry and exit of this procedure. A return address was placed on the stack when `set_stack` was called. Because this code switches stacks, this address is popped from the stack on entry and pushed on the stack before exit. The return stack procedure then restores the stack segment and pointer to what was saved in `set_stack`.

**Listing 9 10: Stack Switch Module STACK.ASM**

```

;STACK.ASM
;Define segment names used by C
;
;_TEXT segment byte public 'CODE'
;_TEXT ends
;CONST segment word public 'CONST'
;CONST ends
;_BSS segment word public 'BSS'
;_BSS ends
;_DATA segment word public 'DATA'
;_DATA ends
;
;GROUP GROUP CONST, _BSS, _DATA
;
;assume CS:_TEXT, DS:GROUP
;
;public _set_stack, _restore_stack

```



```

extrn  _stack_ptr:near      ;our TSR stack
extrn  ss_save:near        ;save foreground SS
extrn  _sp_save:near       ,save foreground SP

_TEXT segment
;****
;void far set_stack(void) -
; save current stack and setup our local stack
;****
_set_stack proc far
;save foreground stack

;we need to get the return values from the stack
;since the current stack will change
    pop ax ;get return offset
    pop bx ;get return segment

;save away foreground process' stack
    mov word ptr _ss_save,ss
    mov word ptr _sp_save,sp

;setup our local stack
    mov ss,word ptr _stack_ptr+2
    mov sp,word ptr _stack_ptr

IFDEF MULTI
    mov bp,sp ;make bp relative to our stack frame
ENDIF
;setup for ret
    push bx
    push ax
    ret
_set_stack endp
;****
;void far restore_stack(void) -
; restore foreground stack, throw ours away
;****
_restore_stack proc far
;we need to get the return values from the stack
;since the current stack will change
    pop cx ;get return offset
    pop bx ;get return segment

;save background stack
    mov word ptr _stack_ptr+2,ss
    mov word ptr _stack_ptr,sp

;restore foreground stack here
    mov ss,word ptr _ss_save
    mov sp,word ptr _sp_save

IFDEF MULTI
    mov bp,sp ;make bp relative to our stack frame
ENDIF
;setup for ret
    push bx
    push cx
    ret
_restore_stack endp
_TEXT ends
_DATA segment
_DATA ends
end

```

of the stack is that it is vital to compile Microsoft C with `_C_` or with a switch that disables `_W_` or turn off stack checking. Otherwise, the C compiler's `chkstk` routine would get hopelessly confused by the new stack.

## DOS Functions for TSRs

Let's see what it takes to discuss TSR programming with DOS functions. Recall that the whole issue is how to make `INT 21h` calls from the resident portion of a TSR. First we present the traditional set of `INT 21h` interrupt DOS, then we show the undocumented `INDDOS` technique.

### MS-DOS Flags

DOS keeps a set of memory called the `InDOS` flag, also known as the "DOS safe" flag. This flag indicates a way to avoid accessing `INT 21h` calls if a semaphore that turns DOS into a serially interruptible system is set. The `INDDOS` flag, chapter 8 shows where `INDDOS` is set and decrements the `InDOS` flag, and `INT 21h` is allowed to enter DOS, that is, make `INT 21h` calls, if the semaphore decrements `InDOS` to 1. As you'll see, there are some major exceptions to this rule, but the basic idea is sound.

Generally, the resident portion of a TSR that uses `INDDOS` checks this flag. If the flag indicates that `INDDOS` does not allow the TSR program, must delay the activation, or at least that portion of the activation, of `INT 21h`, the `INDDOS` `INT 21h` function 34h returns the address of the `InDOS` flag. The `INDDOS` system returns a different `INDDOS` as constant for a particular operating environment; the `INDDOS` system is a TSR also, but is not on our list. You can then store the address at a local variable to be accessed from `INDDOS`.

The `C:\DOS\NDDOS.C` file, step 9 of the `DosDbase` that returns a zero if it is safe to make `INT 21h` calls. If `INDDOS` cannot be interrupted, it returns non-zero. Your application must call `IntInDos` to set the address of the `InDOS` flags during initialization. If you neglect to do so, `DosDbase` always returns non-zero. Of course, `DosDbase` could instead just call `IntInDos` for you at this point.

The `Int28hDbase` function in `step 9` of this book, if the `InDOS` flag is greater than one or the critical error flag is set to zero, and is intended to be used only inside an `INT 28h` loop. Inside `INT 28h`, `INDDOS` -- if it is more than 1, it calls `DOS` via `INDDOS` 1 inside `INT 28h` means it is busy.

### Listing 9-11: DOS Flag Management Module `INDDOS.C`

```
/* INDDOS.C - functions to manage DOS flags */
```

```
#include <stdio.h>
#include <dos.h>

#define GET_INDDOS    0x34
#define GET_CRIT_ERR  0x5006

char far *indos_ptr = 0;
char far *crit_err_ptr = 0;
int  DosBusy(void);
void  InitInDos(void);
*****
Function: Init InDos Pointers
Initalize pointers to InDos Flags
*****
void InitInDos(void)
{
    union REGS regs;
    struct SREGS sregs;
```



## Get/Set PSP

As discussed in Chapter 7, particularly in the section "Unique Process Identifier," each process in MS-DOS has a Program Segment Prefix. You learned that Memory Control Blocks are stamped with the PSP of their owner. In Chapter 8, you saw that this 256-byte area contains, among other things, a default file handle table, Job File Table, for the process, because it is a unique value (though, as chapters 3 and 6 note, this can get complicated in 386/multitasking environments), the segment address of the PSP also acts as a unique process identifier.

At any given moment in an MS-DOS system, there is a current PSP. In the appendix entry for INT 21h Functions 5100h and 510Bh, you can see that the current PSP is kept at offset 10h in both versions of the DOS Swappable Data Area. When DOS receives an INT 21h Function 30h request to open a file, for example, the handle returned to the user is pushed into the HFI of the current PSP.

Well, that's obviously the PSP that belongs to whatever process is calling INT 21h Function 30h, right? Not necessarily; that we are talking about ISR here. The current PSP, unless you somehow change it, belongs to whatever process happens to be running when we pop up. Thus, if the ISR decides to start opening files, it would be using the foreground process's PSP. This could be totally disastrous.

Consider the following example of a TSR that ignores the current PSP when it pops up. If the TSR opens a file handle or allocates memory using MS-DOS basic call subroutines associated with the foreground process, the foreground process is not a state of course, so the entries in HFI are corrupted. When the foreground terminates, all open files are closed and allocated memory segments are freed. This includes those which the TSR thought it was going to be allowed to be associated with the foreground process. Furthermore, if the TSR opens a resource, pops up, loses one PSP, and tries to read it or use it while popped up, over a different SP, because it has no reference to the wrong file!

What the ISR must do when it pops up, somehow, change DOS's current PSP so that it correctly to the ISR's, carry out whatever task the ISR is supposed to perform when it pops up. And then, when leaving back into its dominant state, restore the correct DOS's current value of SP when the ISR popped up.

Fortunately, DOS provides just the functions you need: DOS INT 21h Functions 50h and 51h are the Get PSP and Set PSP functions. MS-DOS 3.x and above, to DOS 3.x and above, documented Function 62h is also available to Get PSP. As with Get Computer System, then thought that Get PSP returns the PSP of whatever program called it. As with Get System, however, it gets DOS's current PSP, not of the SDA. Likewise, Set PSP sets this same DOS's.

In DOS 3.0 and higher, Functions 50h, 51h, and 62h do not use the DOS stacks and are fully resident. They are among the few INT 21h functions that do not pay attention to the E-DOS flag, and therefore consist of an absolute comparison to the DOS flag. But, as noted earlier, to call any of 50h or 51h in DOS 2.x, you must first set the critical error stack. The following pseudocode describes the steps to use Get/Set PSP's from a TSR.

### TSR\_INITIALIZATION:

```
psp_addr =
Get current PSP with function 50h or 62h
(this PSP will be that of the TSR)
```

```
Terminate and stay resident
```

### TSR\_ACTIVATION:

```
fgrnd_psp_addr =
Get current PSP with function 50h or 62h
(since the TSR interrupted the foreground, this address
will be that of the foreground process)
```

```
Set current PSP with function 50h, using psp_addr
Do TSR work
```

Set current PSP with function 50h, using `fgnd_psp_addr`  
Go back to sleep

Listing 9.12 PSP.C contains functions that Get and Set the current PSP, taking into account the various oddities which we have discussed. These are not just simple-minded sugar coating for the equivalent DOS functions. We test for DOS 2.x and set the critical error flag accordingly. Again, you must call `InitInDos` before using these functions.

### Listing 9.12: PSP Management Module PSP.C

```

/* PSP.C */
#include <stdlib.h>
#include <dos.h>
#include "ter.h"
#define GET_PSP_DOS2 0x51
#define GET_PSP_DOS3 0x62
#define SET_PSP 0x50
extern union REGS regs;
/****
Function: GetPSP - returns current PSP
****/
unsigned GetPSP(void)
{
    if (_osmajor == 2)
    {
        if (!crit_err_ptr) /* forgot to call InitInDos */
            return 0; /* gosh, I should just call InitInDos for them */
        *crit_err_ptr = 0xFF; /* force use of proper stack */
        regs.h.ah = GET_PSP_DOS2;
        intdos(&regs,&regs);
        *crit_err_ptr = 0;
    }
    else
    {
        regs.h.ah = GET_PSP_DOS3;
        intdos(&regs,&regs);
    }
    return regs.x.bx;
}
/****
Function: SetPSP - sets current PSP
****/
void SetPSP(unsigned segPSP)
{
    if (!crit_err_ptr) /* forgot to call InitInDos */
        return; /* should call InitInDOS for them! */
    if (_osmajor == 2)
        *crit_err_ptr = 0xFF; /* force use of correct stack if DOS 2.x */
    regs.h.ah = SET_PSP;
    regs.x.bx = segPSP; /* pass segment value to set */
    intdos(&regs,&regs);
    if (_osmajor == 2)
        *crit_err_ptr = 0; /* restore crit error flag */
}

```

### Extended Error Information

Consider the following scenario: The foreground program has performed a DOS function that failed. Because of this, DOS has stored extended error information. Normally, the foreground

program can access the extended information at this point. The ISR becoming active, however, delays the access of the extended error information. Now the ISR has control and could possibly perform I/O functions that overwrite the existing extended error information, making it invalid for the interrupted (foreground) program.

Don't despair though. DOS has this problem well in hand. In DOS 3.0 and higher, documented Function 57h is available to query the extended error information. A TSR must save this information prior to activation and reset it at exit. DOS Function 5700h allows the extended error information to be set directly to function 5700h, point DS:DX, not DS:SI as claimed in the *MS-DOS Programmer's Reference* for DOS 5. To a table containing the contents of the registers when an error occurred. The register values for this table can be retrieved using Function 59h.

The C functions in Listing 9.13 can get the extended error information on ISR activation and reset the extended error information on exit.

### Listing 9.13. Extended Error Management Module EXTERR.C

```

/* EXTERR.C - extended error saving and restoring */
#include <stdio.h>
#include <dos.h>

#define GET_EXTERR    0x59
#define SET_EXTERR    0x5d0a

#pragma pack(1)
struct ExtErr
{
    unsigned int  errax,  errbx,  errcx,  errdx,  errs1,  errds,
    unsigned int  errds,  errcs,
    unsigned int  reserved,  userID,  programID,
};

void GetExtErr(struct ExtErr * ErrInfo);
void SetExtErr(struct ExtErr near * ErrInfo);

extern union REGS regs;
extern struct SREGS sregs;

/*****
Function: GetExtErr
get extended error information
*****/
void GetExtErr(struct ExtErr * ErrInfo)
{
    if (_osmajor >= 3) /* only for DOS 3 and above */
    {
        regs.h.ah = GET_EXTERR,
        regs.x.bx = 0; /* must be zero */
        intdosx(&regs, &regs, &sregs);
        ErrInfo->errax = regs.x.ax,
        ErrInfo->errbx = regs.x.bx,
        ErrInfo->errcx = regs.x.cx,
        ErrInfo->errdx = regs.x.dx,
        ErrInfo->errs1 = regs.x.si,
        ErrInfo->errds = regs.x.di,
        ErrInfo->errds = sregs.ds,
        ErrInfo->errcs = sregs.es,
        ErrInfo->userID = 0, /* zero per DOS 3 Tech Ref */
        ErrInfo->programID = 0,
    }
}

/*****
Function: SetExtErr

```

```

set extended error information
****/
void SetExtErr(struct ExtErr near * ErrInfo)
{
    if (_osmajor >= 3) /* only for DOS 3 and above */
    {
        regs.x.ax = SET_EXTERR,
        regs.x.bx = 0; /* must be zero */
        segread(&regs), /* put address of err info in DS:DI */
        regs.x.di = (int) ErrInfo,
        intdosx(&regs,&regs,&regs);
    }
}

```

### Extended Break Information

At times it is necessary to tell the operating system "Hey, stop what you're doing and go back to the command prompt!" This is better known as Ctrl-C or Ctrl-Break.

MS-DOS, by default, checks the interrupt keys for these break characters only when you use the CP/M emulation. Operations INT 21h: AH=01h (page 04), BREAK ON. However, DOS checks for the break key when processing a most all INT 21h functions, see the INT 21h dispatch above. It is unexpected. While at the COMMAND prompt, capturing BREAK ON does extended break checking and BREAK OFF turns off the extra check.

Now, what does this have to do with a TSR?

If extended break checking is on and you press a break key just before the TSR becomes active in MS-DOS, instead of ending the TSR, it's! This is a small window of disaster, but it would probably happen just as you are likely to save that area program you were editing.

To avoid this possibility, you must query and save the current break status. Then, at the first DOS interrupt or call within the TSR activation section, must set to turn extended break checking off. On TSR exit, restore the original break status.

The one problem to be solved on concrete DOS 2.x, where getting or setting the extended break status causes a break key's pending, and extended break checking is ON. This is corrected in DOS 3.0 and later. About any, there is no clean solution for DOS 2.x. To avoid problems with the TSR examples in this chapter, extended break checking must be OFF for DOS 2.x.

Listing 9-14 shows BREAK.C, containing 3 functions to get and set the Extended break status.

### Listing 9-14: Extended Break Management Module BREAK.C

```

/* BREAK.C */
#include <stdlib.h>
#include <dos.h>
#include <far.h>
#define GET_SET_BREAK 0x33
#define GET_BREAK 0
#define SET_BREAK 1
extern union REGS regs;

*****/
Function: GetBreak - returns current Break Status 0 = OFF, 1 = ON
*****/
int GetBreak(void)
{
    if (_osmajor != 2)
    {
        regs.h.ah = GET_SET_BREAK;
        regs.h.al = GET_BREAK;
    }
}

```





Thus our ISR has acquired another wrinkle. In addition to checking the IDIOS and clearing error flags, saving and restoring IDIOS's current PSP, and saving and restoring the extended I/O information, you must now also hook INI 28h (see Web page) ever since the IDIOS ISR interface was a model of clarity. It is unlikely that anyone at Microsoft sat down and tried to design "nice" interfaces for ISR programming. Instead they put in what they needed to write their own ISRs. The end result is an interface that looks like something someone would design only for their own use. On the other hand, the IDIOS ISR interface, if we can call it that, benefits from the fact that its designers used it themselves. They ran it to the same problems you run into with other ISRs, so they put it solutions.

In error handlers for INI 28h should pass control to the previous INI 28h routine, which is complete. Generally they should not hog the INI 28h interrupt by executing large amounts of code. ISRs that solicit user input should not only hook INI 28h but also periodically invoke INI 28h to their input loop. This gives other ISRs a chance to be scheduled. For our generic ISR, we use INI 28h to detect whether the user had earlier pressed the hotkey at a time when we could "pop up" to the MUI 11 ISR program at the end of this chapter. INI 28h's whole idea is that we can use the time slices we get while the system is sitting at the COMMAND prompt.

The IDIOS C mode's `carrier` variable contains the function called `INI28h` below, which returns zero if it is safe to access IDIOS during an INI 28h. The `INI00S` flag is never zero during an INI 28h, so you might think that you don't need to ever check `INI00S` during an INI 28h. However, `INI00S` could be greater than one, in which case you still need a pop-up.

The INI 28h handler appears in the main ISR module, `ISREXAMP.C`, to which we now turn

## Inside the Generic TSR

The main module for our generic ISR, `ISREXAMP.C`, initializes the program, then it is pop-up routine calls your application to do several interrupt handlers. The remainder of the interrupt handlers are written in assembly language. For reasons noted earlier and are found in `ISREXAMP.SAT` and `ISREXAMP.C`, we have to use the modules we use already exist: `IDIOS.C`, `PSP.C`, `EXTERR.C`, and `STACK.ASM`.

But the big picture directly into a 850 lines of code belonging to `ISREXAMP.C`, the 251 lines that compose `ISREXAMP.ASM` will start off with a pseudocode explanation. By coloring pseudocode keywords to use of the keyword `ON`, borrowed from `ASM`, which is only borrowed from `AT&T`. A structure class `ON` (MIL) records the code that is scheduled for the program's execution from inside the program. It is a very convenient reader with only a C using the interrupt keyword, and earlier a C work using it. This example does several things with `ON`. The `ON` keyword is used in the following pseudocode's particular expressions of what happens in our ISR.

The first while discussion assumes that the ISR is going to access the disk during its pop-up phase and that the ISR does not use the `IDIOSWAP` interface, which we've mentioned, but not yet discussed in detail.

The initialization of the generic ISR looks something like this:

```
INIT ; main() in ISREXAMP.C
  IF they want to deinstall
    CALL deinstall(C)
  ELSE IF TSR not already installed
    MALLOC stack
    GETVECT TIMER(8), KEY(9), DISK(13h), IDLE(2Bh), MULTIPLEX(2Fh)
    ; 2Fh is for communication with already resident copy of TSR
  ; (install check, deinstall)
  SETVECT TIMER, KEY, DISK, IDLE, MULTIPLEX
```

```
RELEASE environment
RELEASE unused heap
TSR
```

If there are two presses here, except perhaps the fact that we are somehow using INT 2Fh to communicate between an already resident copy of the TSR program and a second copy that rather than go TSR deinstalls to resident copy. If INT 1Fh has already been made resident, it can be deinstalled by typing INT 1Fh at the COMMAND prompt. If you are starting rather than deinstalling, then the TSR first checks to see if the TSR is already installed, as well as INT 2Fh mode status. If it isn't installed, then the INT routine completes the program, as well as clearing the resident, and five interrupt handlers have been installed.

Before examining the interrupt handlers, let us create a few subroutines that the interrupt handlers will use to communicate among themselves:

```
FLAG wanted_pop_up ; wanted to pop up earlier, but DOS was busy
FLAG disk_unsafe ; INT 13h in use?
FLAG idle_int ; is INT 2Fh in progress?
```

EDOS is not included here because this flag is maintained by EDOS itself, and is not located inside the program. We use our DosBusy routine to check EDOS.

The first interrupt handler to examine is the one that handles keyboard events. Because we also have an INT 9 handler, each time the user presses any key, a copy of the address of something like the following gets executed:

```
ON KEY ; new_int9() in TSREXAMP.C
IF we are not already running AND
IF it's our hotkey AND
IF NOT disk_unsafe THEN
CALL POP_UP
ELSE
; we can't pop up now, so just set flag indicating that
; we WANT to pop up at the next available moment
wanted_pop_up = TRUE
JMP previous KEY handler
ELSE
; not our hotkey - chain to next handler
JMP previous KEY handler
ELSE
; we're already running - let key be processed normally
JMP previous KEY handler
```

By the simplest scenario, the user presses the hotkey at a time when INT 13h isn't in use. Our key board handler then calls the POPUP routine:

```
POP_UP ; isr_function() in TSREXAMP.C
CALL set_stack() ; switch to our own stack
IF DosBusy() AND NOT idle_int
wanted_pop_up = TRUE
ELSE
; we really can POP UP now!
GETVEC? CTRL-BREAK(F0h), CTRL-C(1Ch), CR(10h,20h)
SETVEC? CTRL-BREAK, CTRL-C, CR(10h)
current_PSP = GetPSP()
CALL SetPSP(1SR_PSP) ; 1SR_PSP and 1SR_DTA were set in init
current_DTA = GetDTA()
CALL SetDTA(1SR_DTA)
save_err = GetExtErr()
eat keys
CALL application()
CALL SetExtErr(save_err)
CALL SetDTA(current_DTA)
CALL SetPSP(current_PSP)
```

```
SETVECT CTRL BREAK, CTRL C, CRITERM, REVERT
```

```
ON CTRL-BREAK DO NOTHING
```

```
ON CTRL-C DO NOTHING
```

```
ON CRITERM, new_int24() in TSREXAMP C
RETURN FAILURE
```

Continuing with the simplest scenario let's say that DosBios returns IAMS. The program then proceeds to install three short-term interrupt handlers. The Ctrl-Break and Ctrl-C handlers merely discard these events. A more sophisticated ISR might do something fancy with them. The Criteria Error handler merely returns failure. The key point is that the pop-up portion of the ISR must install its own handlers for these events, not whatever handler the foreground process happens to have installed at the time. Next, the ISR swaps its own Disk Transfer Acc and PSP with that of the foreground process, and saves the extended error information it received, if any. It catches whatever keys are lurking in the keyboard buffer, and finally calls the application, which, as you know, does something useful like providing a notepad during a mode-*n* application taken from *Programs*. When the application finishes, the ISR puts everything back the way it found it.

That was the simplest scenario. Say the user has pressed the hotkey, but the ISR can't pop up. Either INT18 has been disabled, or BIOS is calling busy—that is, INT0DS is set and BIOS won't issue an INT28h idle interrupt. In this case, either the keyboard handler or the pop-up routine sets the wanted\_popup flag, and then calls keyboard\_handler in the case of the keyboard handler, with an IRQ1.

So all the ISR has to do is set the wanted\_popup flag. How is this going to help the ISR pop up?

Remember the INT18 timer tick handler the ISR installed. At each timer tick, about 18.2 times a second, unless someone has reprogrammed the chip that generates these interrupts, our TIMER routine gets invoked. Its job is to check the wanted\_popup flag.

```
ON TIMER ; new_int8() in TSREXAMP C
CALL previous TIMER handler
IF NOT tsr_active
IF wanted_popup
IF NOT DosBusy()
IF NOT disk_unsafe
wanted_popup = FALSE
CALL POP UP
```

Once the wanted\_popup flag has been set, the TIMER routine checks 18.2 times a second to see if it's now safe to pop up. It does so until it is safe to pop up, at which time the flag is nulled out.

One thing not shown in pseudocode, but appearing in the genuine code in ISR1.VAMP.C and ISR1.IDLE.VAMP.C, is that all valid error entering is via INT18 or INT19. The ISR cares to the previous error. In addition to giving the previous interrupt handler an opportunity to do its thing, it also receives the previous saved error and the end of interrupt (EOI) command (IOPLT=0x8259) interrupt, plus control of this system, which you won't find there, since obligatory call to out(0x20,0x20) sprinkled throughout the code.

In addition to timer ticks, you can also use the idle interrupt as a trigger for servicing a wanted\_popup request. Note also that the IDLE handler increments and decrements the idle int flag, which is checked at entry to the POPUP routine.

```
ON IDLE ; new_int28() in TSREXAMP C
INCR idle int
IF wanted_popup
IF NOT int28DosBusy()
IF NOT tsr_active
IF NOT disk_unsafe
POP UP!
```

```
DECR idle_int
CALL previous IDLE handler
```

Finally, how does the important disk unsafe flag stay updated? There is unfortunately no flag in the BIOS that we can get a far pointer to as we did with IrPCBS, so we hook INT 13h to create our own disk unsafe semaphore.

```
ON DISK ; new_int13() in TSRUTIL.ASM
INCR disk_unsafe
CALL previous DISK handler
DECR disk_unsafe
```

That's about all there is to the generic ISR. Note how decentralized the code for a ISR is. Rather than have one top-level routine that calls various sub-routines, there is instead a collection of independent handlers that get called due to some event taking place outside the program. The system has a "top" and instead consists of these asynchronously invoked agents. Much is made of event driven programming in environments like Windows, but in reality it's not much different from what we're doing here.

Having taken this walk through the pseudocode, you should now be able to fully understand the actual low C source code in SREXAMP.C (Listing 9-15). However, many of the variable and function names are different from our pseudocode, and we haven't yet explained the sections which are conditionally compiled with #ifdef DOS5WAP.

SREXAMP.C includes the rather interesting, but necessary, ISR.H, which contains typedefs and function prototypes for all the modules that make up the generic ISR. Listing 9-16 shows ISR.H.

### Listing 9-15: TSREXAMP.C

```
/*
TSREXAMP.C
by Raymond J. Nichols
with revisions by Andrew Schulman
Second Edition - includes MS-DOS Task Manager/Windows Support
*/

#include <stddef.h>
#include <stdlib.h>
#include <stdio.h>
#include <conio.h>
#include <dos.h>
#include <bios.h>
#include <memory.h>
#include "tsr.h"

#define STACK_SIZE 8192 /* must be 16 byte boundary */
#define SET_DTA 0x1a /* SET Disk Transfer Address */
#define GET_DTA 0x2f /* GET Disk Transfer Address */
#define DOS_EXIT 0x4c /* DOS zero name (exit) */
#define KEYBOARD_PORT 0x60 /* KEYBOARD Data Port */
#define PSP_TERMINATE 0x0a /* Termination addr. in our PSP */
#define PSP_PARENT_PSP 0x16 /* Parent's PSP from our PSP */
#define PSP_ENV_ADDR 0x2c /* environment address from PSP */
#define HOT_KEY 32 /* Hot key a long with ALT (B)*/
#define RIGHT_SHIFT 1
#define LEFT_SHIFT 2
#define CTRL_KEY 4
#define ALT_KEY 8
#define MULTIPLEX_ID 0xcd
#define INSTALL_CHECK 0x00
```

```

#define INSTALLED      0xFF
#define DEINSTALL     0x01
#define SESSION_ACTIVE 0x03

#define PARAGRAPHS(x) ((FP_OFF(x) + 15) >> 4)

unsigned char multiplex_id = MULTIPLEX_ID,
char far *stack_ptr;      /* pointer to TSR stack */
unsigned ss_save,         /* slot for stack segment register */
unsigned sp_save,        /* slot for stack pointer register */
volatile int int_2B_in_progress = 0; /* true if INT 2Bh in progress */
volatile int unsafe_flag = 0; /* true if INT 13h in progress */
unsigned keycode;
char buf[20],            /* work buffer */
unsigned long TerminateAddr, /* used during de-install */
union REGS regs;        /* register work structures */
struct SREGS sregs,
int hot_key,           /* keycode for activation */
int shift_key,         /* shift status bits (alt, ctrl..) */
int user_key_set = 0;

/* Save areas for old interrupt pointers */
INTVECT old_int8, old_int9, old_int10, old_int13,
INTVECT old_int2B, old_int2f;

#ifdef DOS_SWAP
extern int dos_critical; /* used by DOSSWAP.C */
INTVECT old_int2a,
void interrupt far new_int2a(INTERRUPT_REGS);
#endif

extern int enhanced_windows,
extern int switcher_critical;

/* Global Data block that will be maintained in Instance Memory
That way each session in a Task Manager will have unique Global
Data. */

struct
{
    volatile int tsr_already_active, /* true if TSR active */
    volatile int popup_while_dos_busy, /* true if hotkey hit while dos busy */
    int breakState, /* status of MS-DOS break checking */
    INTVECT old_int1b, old_int23, old_int24, /* PSP of process we've interrupted */
    unsigned foreground_psp; /* DTA of process we've interrupted */
    unsigned foreground_dta_seg;
    unsigned foreground_dta_off;
    struct ExtErr ErrInfo, /* save area for extended error info */
} TerGlb;

int g_obs.tsr_active = 0, /* flag to indicate a session is in TSR */

/* PROTOTYPES FOR THIS MODULE */
void interrupt far new_int8(INTERRUPT_REGS),
void interrupt far new_int9(INTERRUPT_REGS),
extern void interrupt far new_int13(void), /* in TSRUTIL.ASM */
void interrupt far new_int1b(INTERRUPT_REGS),
void interrupt far new_int23(INTERRUPT_REGS),
void interrupt far new_int24(INTERRUPT_REGS),
void interrupt far new_int2B(INTERRUPT_REGS),
void interrupt far new_int2f(INTERRUPT_REGS),
void tsr_function(void),
void tsr_exit(void),
void usage(char *);
int Dk, intVect(int Vect, INTVECT NewInt, INTVECT OldInt),
void parse_cmd_line(int argc, char *argv[]);
void main(int argc, char *argv[]);
*****
* TIMER INTERRUPT HANDLER

```

```

*****;
void interrupt far new_int8(INTERRUPT_REGS r)
{
    (*old_int8)(); /* process timer tic */
#ifdef DOS_SWAP
    if (!TsrGlb.tsr_already_active && TsrGlb.popup_while_dos_busy &&
        'dos_critical && 'unsafe_flag)
#else
    if (!TsrGlb.tsr_already_active && TsrGlb.popup_while_dos_busy &&
        'DcsBusy() && 'unsafe_flag)
#endif
    {
        /* If we're not running enhanced mode windows, set flag so that
           task swap will not occur. If we're not running Windows,
           enhanced defaults to true, so there is no effect */
        if ('enhanced_windows)
            switcher_critical++, /* don't allow task switch */
        TsrGlb.popup_while_dos_busy = 0,
        TsrGlb.tsr_already_active = 1,
        global_tsr_active = 1,
        _enable(); /* turn interrupts back on */
        tsr_function(),
        TsrGlb.tsr_already_active = 0;
        global_tsr_active = 0;
        if ('enhanced_windows)
            switcher_critical--, /* task switch ok */
    }
}

/*****
 * KEYBOARD INTERRUPT HANDLER
 *****/
void interrupt far new_int9(INTERRUPT_REGS r)
{
    if (!TsrGlb.tsr_already_active)
    {
        if ((keycode = inp(KEYBOARD_POR)) != 'not key)
            _chain_intr(oid_int9),
            if (('b'os_keybrd( KEYBRD_SHIFTSTATUS) & shift_key) == shift_key)
            {
#ifdef USES_DISK
                if ('unsafe_flag)
                {
                    TsrGlb.popup_while_dos_busy = 0,
                    TsrGlb.tsr_already_active = 1,
                    global_tsr_active = 1,
                    (*old_int9)(); /* send key to old int routine */
                    tsr_function(),
                    TsrGlb.tsr_already_active = 0,
                    global_tsr_active = 0;
                }
#ifdef USES_DISK
            }
            else
            {
                TsrGlb.popup_while_dos_busy = 1,
                _chain_intr(oid_int9),
            }
        }
    }
    else
        _chain_intr(oid_int9);
}

```

```

    }
    else
        _chain_intr(oid_int9);
}
/*****
 * CTRL-BREAK INTERRUPT HANDLER
 *****/
void interrupt far new_int1b(INTERUPT_REGS r) { /* do nothing */ }
/*****
 * CTRL-C INTERRUPT HANDLER
 *****/
void interrupt far new_int23(INTERUPT_REGS r) { /* do nothing */ }
/*****
 * CRITICAL ERROR INTERRUPT HANDLER
 *****/
void interrupt far new_int24(INTERUPT_REGS r)
{
    if (_osmajor >= 3)
        r.ex = 3; /* fail dos function */
    else
        r.ex = 0;
}
/*****
 * DOS IDLE INTERRUPT HANDLER
 *****/
void interrupt far new_int28(INTERUPT_REGS r)
{
    int_28_in_progress++;
#ifdef DOS_SWAP
    if (!TsrGlb.popup_while_dos_busy && 'dos_critical'
        && 'TsrGlb tsr_already_active && 'unsafe_flag')
#else
    if (!TsrGlb.popup_while_dos_busy && ('int28dosbusy()
        && 'TsrGlb tsr_already_active && 'unsafe_flag')
#endif
    {
        TsrGlb.tsr_already_active = 1;
        global_tsr_active = 1;
        tsr_function();
        TsrGlb.tsr_already_active = 0;
        global_tsr_active = 0;
    }
    int_28_in_progress--;
    _chain_intr(oid_int28);
}
#ifdef DOS_SWAP
/*****
 * DOS INTERNAL INTERRUPT HANDLER
 *****/
void interrupt far new_int2a(INTERUPT_REGS r)
{
    switch (r.ex & 0xff00)
    {
        case 0x8000: /* start critical section */
            dos_critical++;
            break;
        case 0xB100: /* end critical section */
        case 0xB200: /* end critical section */
            if (dos_critical) /* don't go negative */
                dos_critical--;
            break;
    }
}

```

```

    default,
        break;
}
chain_intr(old_int2a);
}
#endif
/*****
 * DOS MULTIPLEX INTERRUPT HANDLER
 *****/
void interrupt far new_int2f(INTERRUPT_REGS r)
{
    if (process_switcher_int2f(&r)) /* check if task switcher function */
    {
        unsigned ah = r.ax >> 8;
        unsigned al = r.ax & 0xFF;
        if (ah != multiplex_id)
        {
            if (al == INSTALL_CHECK)
                r.ax |= INSTALLED;
            else if (al == DEINSTALL)
            {
                if (global_isr_active)
                    /* a session is active, don't allow de-install */
                {
                    r.ax &= 0x100,
                    /*let caller know we're still there
                    r.ax |= SESSION_ACTIVE;
                    }
                else
                {
                    // because of stack swap, pass arg in static variable
                    TerminateAddr = ((long)r.bx << 16) + r.dx,
                    if (!TsrGlb_isr_already_active)
                        /* don't exit if we're active */
                    {
                        _enable(); /* STI */
                        Tsr_exit();
                        // If we got here, we weren't able to unlink
                        // let caller know we're still there
                        r.ax = 0xFFFF;
                        // MSC 6.0 /Ox optimizes out the above
                        // get it back by using the value in ax
                        TsrGlb_isr_already_active = ~r.ax,
                        //set to 1 to prevent any more action
                    }
                }
            }
        }
        else
            chain_intr(old_int2f);
    }
}
/*****
 * ISR ACTIVE SECTION
 *****/
void tsr_function()
{
    switcher_critical++;
    windows_begin_critical(),
    set_stack(),
    windows_end_crit call(),
    switcher_critical ;
}

```



```

#ifdef DOS_SWAP
    if (SaveDosSwapf) && !int_28_in_progress)
#else
    if (DosBusyf) && !int_28_in_progress)
#endif
    /* set flag: next INT 8,28 activates us */
    TsrGlb.popup_while_dos_busy = 1;
else
{
    TsrGlb.popup_while_dos_busy = 0;
    /* save current extended break status */
    TsrGlb.breakState = GetBreak();
    /* set extended break status to be off */
    SetBreak(0),
}

#ifdef DOS_SWAP
    /* Get Extended Error Information */
    GetExtErr(&TsrGlb.ErrInfo);
#endif

    /* save old interrupt-CTRL BREAK, CTRL-C and CRIT ERROR */
    TsrGlb.old_int1b = _dos_getvect(0x1b);
    TsrGlb.old_int23 = _dos_getvect(0x23);
    TsrGlb.old_int24 = _dos_getvect(0x24);
    /* set our interrupts functions */
    _dos_setvect(0x1b, new_int1b);
    _dos_setvect(0x23, new_int23);
    _dos_setvect(0x24, new_int24);

    /* not needed for DOS_SWAP, but can be used by application */
    TsrGlb.foreground_psp = GetPSP();

    SetPSP(_psp), // _psp in STDLIB.H

#ifdef DOS_SWAP
    /* get foreground DTA */
    regs.h.ah = GET_DTA,
    intdosx(&regs, &regs, &regs);
    TsrGlb.foreground_dta_seg = sregs.es,
    TsrGlb.foreground_dta_off = regs.x.bx,
#endif

    /* set up our DTA */
    regs.h.ah = SET_DTA,
    regs.x.dx = 0x80; /* use default in PSP area */
    sregs.ds = _psp,
    intdosx(&regs, &regs, &regs),

    /* suck up key(s) in buffer */
    while (!_bios_keybrd(KEYBRD_READY))
        _bios_keybrd(KEYBRD_READ),

    /* your code goes here */
    applicationf);

#ifdef DOS_SWAP
    /* put back original INTS */
    _dos_setvect(0x1b, TsrGlb.old_int1b);
    _dos_setvect(0x23, TsrGlb.old_int23);
    _dos_setvect(0x24, TsrGlb.old_int24);
    /* put back extended error information */
    SetExtErr(&TsrGlb.ErrInfo),

```

```

RestoreDosSwap();

#else

/* put back original DTA */
regs.h.eh = SET_DTA;
regs.x.dx = TsrGlb.foreground_dta_off;
sregs.ds = TsrGlb.foreground_dta_seg;
intdos(&regs, &regs, &sregs);
/* put back original PSP */
SetPSP(TsrGlb.foreground_psp);
/* put back original INTS */
_dos_setvect(Dx1b, TsrGlb.old_int1b),
_dos_setvect(Dx23, TsrGlb.old_int23),
_dos_setvect(Dx24, TsrGlb.old_int24);
/* put back extended error information */
SetExtErr(&TsrGlb.ErrInfo);

#endif

/* put back extended break checking the way it was */
SetBreak(TsrGlb.breakState),

)
switcher_critical++;
windows_begin_critical();
restore_stack();
windows_end_critical();
switcher_critical--;
}
// only restores 0:dint if someone hasn't grabbed away Vect
int UnlinkVect(int Vect, INTVECT Newint, INTVECT Oldint)
{
    if (Newint == _dos_getvect(Vect))
        ( _dos_setvect(Vect, Oldint), return 0, )
    return 1;
}

void tsr_exit(void)
{
    set_stack();
    /* put interrupts back the way they were, if possible */
    if (!UnlinkVect(8, new_int8, old_int8) |
        !UnlinkVect(9, new_int9, old_int9) | // Do not use I, we
        !UnlinkVect(0x28, new_int28, old_int28) | // DON'T want early out
        !UnlinkVect(0x13, new_int13, old_int13) |
#ifdef DOS_SWAP
        !UnlinkVect(0x2a, new_int2a, old_int2a) |
#endif
        !UnlinkVect(0x2f, new_int2f, old_int2f) )
    {
        // Set parent PSP, &red n: our own PSP, to the current PSP
        *(int far * ((long) _psp << 16) + PSP_PARENT_PSP) = GetPSP(),
        // Set terminate address in our PSP
        *(long far * ((long) _psp << 16) + PSP_TERMINATE) = TerminateAddr,
        SetPSP(_psp); /* set psp to be ours */
        bdos(DOS_EXIT, 0, 0); /* exit program */
    }
    restore_stack(),
}

```

```

void usage(char *prognam)
{
    fputs("Usage: ", stdout);
    puts(prognam);
    puts( [ d deinstall] [ k key shift-keys] [ f multiplex id] );
    puts(" Valid Multiplex id");
    puts( 00 through 15 specifies a unique INT 2fh ID");
    puts( Valid shift-keys is any combination of-");
    puts(" 1 = Right Shift");
    puts(" 2 = Left Shift");
    puts(" 4 = CTRL");
    puts(" 8 = ALT");
    exit(1);
}

void do_deinstall(char *prognam)
{
    fputs(prognam, stdout);
    switch (deinstall())
    {
        case 1
            puts(" was not installed"); break;
        case 2
            puts(" deinstalled"); break;
        case SESSION_ACTIVE
            puts(" TSR active in another session. "
                "TSR was not deinstalled"); break;
        default
            puts(" deactivated but not removed"); break;
    }
    exit(0);
}

int set_shift_key(unsigned sh)
{
    /* figure out, report on shift statuses */
    /* make sure shift key < 0x10 and non-zero */
    if ((!(shift_key - sh) < 0x10) && shift_key)
    {
        printf("Activation is: %s\n",
            shift_key & RIGHT_SHIFT ? "RIGHT " : "",
            shift_key & LEFT_SHIFT ? "LEFT " : "",
            shift_key & CTRL_KEY ? "CTRL " : "",
            shift_key & ALT_KEY ? "ALT " : "",
            "hot key");
        return 1;
    }
    else /* error, bad param */
    {
        puts("Invalid Shift-Status");
        return 0;
    }
}

void parse_cmd_line(int argc, char *argv[])
{
    int i;
    int tmp;
    for (i = 1, i < argc, i++) /* for each cmdline arg */
        if (argv[i][0] == '-' || (argv[i][0] == '/')
            switch(toupper(argv[i][1]))
            {
                case 'D':
                    do_deinstall(argv[i]);
                    break;
                case 'K': /* set pop up key sequence */

```

```

user_key_set = 1;
i++; /* bump to next argument */
if ((hot_key = atoi(argv[i])) != 0)
{
    i++; /* bump to next argument */
    if ('set_shift_key(atoi(argv[i]))
        usage(argv[0]);
}
else
    usage(argv[0]);
break;
case 'F': /* set multiplex ID */
i++; /* bump to next argument */
if ((tmp = atoi(argv[i])) < 0x10)
    multiplex_id += tmp; /* range of C0-CF */
else
    usage(argv[0]);
break;
default: /* invalid argument */
    usage(argv[0]);
} /* end switch */
else
    usage(argv[0]);
}

```

```
void FAIL(char *s) { puts(s); exit(1); }
```

```
void main(int argc, char *argv[])
```

```

{
    union REGS regs,
    struct SREGS sregs,
    unsigned far *fp;
    unsigned memtop, dummy;
    int ret_value;

    /* initialize flags */
    *argLib.popup_while_dos_busy = 0;
    *argLib.tar_already_active = 0;
    InitDos();

    /* the following must be called before adding instance data */
    init_switcher_structures();
    parse_cmd_line(argc, argv);
    /* check if TSR already installed? */
    regs.h.ah = multiplex_id;
    regs.h.al = INSTALL_CHECK;
    int86(0x2f, &regs, &regs);
    if (regs.h.al == INSTALLED)
    {
        puts("TSR already installed");
        fputs(argv[0], stdout); puts(" - D de installis");
        exit(1);
    }
    if ('user_key_set')
    {
        puts("Press ALT-D to activate TSR ");
        printf("Multiplex ID = 0x %n", multiplex_id);
        hot_key = HOT_KEY;
        shift_key = ALT_KEY;
    }
}

#ifdef DOS_SWAP
if ((ret_value = InitDosSwap()) != 0.
{

```

```

puts("Error initializing DOS Swappable Data Area");
switch (ret_value)
{
    case 1:
        puts("No Available Memory"), break;
    case 2:
        puts("Too Many Swap Blocks"), break;
    case 3:
        puts("Unsupported DOS"); break;
    case 4:
        puts("Unable to add swap blocks to instance memory"),
        break;
    default:
        puts("Failure Return from GetDosSwap Call"), break;
}
exit(1);
}
#endif
/* WALLOC a stack for our TSR section */
if ((stack_ptr = malloc(STACK_SIZE)) == NULL)
    FAIL("Unable to allocate stack");
if (add_instance_block(stack_ptr, STACK_SIZE))
    FAIL("Unable to add stack to instance data");
stack_ptr += STACK_SIZE;
if (add_instance_block(&tsrGlob, sizeof(TsrGlob)))
    FAIL("Unable to add TSR global data to instance data");
if (add_instance_block(&ss_save, sizeof(ss_save))
    FAIL("Unable to add ss_save global data to instance data");
if (add_instance_block(&sp_save, sizeof(sp_save))
    FAIL("Unable to add sp_save global data to instance data");
if (add_instance_block(&stack_ptr, sizeof(stack_ptr))
    FAIL("Unable to add stack_ptr global data to instance data");
if (add_instance_block(&regs, sizeof(regs))
    FAIL("Unable to add regs global data to instance data");
if (add_instance_block(&srregs, sizeof(srregs))
    FAIL("Unable to add srregs global data to instance data");

/* check if windows running */
check_windows_running();

/* check if task switcher running */
check_switcher_running();

/* get interrupt vector */
o_d_int8 = _dos_getvect(8); /* timer interrupt */
o_d_int9 = _dos_getvect(9); /* keyboard interrupt */
o_d_int13 = _dos_getvect(0x13); /* disk intr. in TSRUTIL.ASM */
o_d_int28 = _dos_getvect(0x28); /* dos idle */
o_d_int2f = _dos_getvect(0x2f); /* multiplex int */

#ifdef DOS_SWAP
o_d_int2a = _dos_getvect(0x2a); /* dos internal int */
#endif

init_intr(), /* initialize int routines in TSRUTIL.ASM */

/* set interrupts to our routines */
_dos_setvect(8, new_int8),
_dos_setvect(9, new_int9),
_dos_setvect(0x13, new_int13), /* in TSRUTIL.ASM */
_dos_setvect(0x28, new_int28),
_dos_setvect(0x2f, new_int2f),

#ifdef DOS_SWAP
_dos_setvect(0x2a, new_int2a),
#endif
#endif

```

```

/* release environment back to MS-DOS */
FP_SEG(fp) = _psp;
FP_OFF(fp) = PSP_ENV_ADDR;
dos_freemem(*fp);

/* release unused heap to MS-DOS */
/* All MALLOCs for TSR section must be done in TSR_INIT() */
/* calculate top of memory, and go TSR */
segread(&regs);
memtop = sregs.es + PARAGRAPHS(stack_ptr) - _psp;
_dos_keep(0, memtop);
)

```

### Listing 9-16. TSR.H

```

/* TSR Prototype file and common variables */
#ifdef cplusplus
extern "C" {
#endif

#define INTERRUPT void interrupt far

typedef struct {
#ifdef TURBOC
    unsigned bp, di, si, ds, es, dx, cx, bx, ax;
#else
    unsigned cu, ds;
    unsigned di, si, bp, sp, bx, dx, cx, ax; /* PUSHA */
#endif
} INTERRUPT_REGS;

typedef void (*interrupt far *INTVECT)();

/* Prototypes for functions in INDOS.C */
int  BusBusy(void);
int  Int2BdsBusy(void);
void InitInDos(void);

/* Prototypes for functions in PSP.C */
unsigned GetPSP(void);
void SetPSP(unsigned segPSP);

/* Prototypes for functions in TSRUTIL.ASM */
int  far  deinstall(void);
void far  init_intr(void);
void far  idle_int_chain(void);

void interrupt far  new_int10(void);
void interrupt far  new_int13(void);
void interrupt far  new_int25(void);
void interrupt far  new_int26(void);

void far  timer_int_chain(void);

/* Prototypes for functions in STACK.ASM */
void far  get_stack(void);
void far  restore_stack(void);

/* Prototypes for functions in EXTER.C */
void GetExtErr(struct ExtErr * ErrInfo);
void SetExtErr(struct ExtErr * ErrInfo);

struct ExtErr
{
    unsigned int  errax, errbx, errcx, errdx, errs1, errd1;
    unsigned int  errds, erres;
    unsigned int  reserved, userID, programID;
};

/* Prototypes for functions in BOSSWAP.C */

```

```

int  InitDosSwap(void),
int  SaveDosSwap(void),
void RestoreDosSwap(void);
/* Pointer defined in INDOS.C */
extern char far *  indos_ptr,
extern char far *  crit_err_ptr;
/* Prototypes for functions in BREAK.C */
int  GetBreak(void);
void SetBreak(int  breakStatus),
/* Prototypes and global vars defined in SWITCHER.C */
extern int  enhanced_windows,
extern int  switcher_critical,
int  process_switcher_int2f(INTERRUPT_REGS far *r),
void  InitSwitcherStructures(void),
int  add_instance_block(void far *ptr, unsigned int size),
void  windows_begin_critical(),
void  windows_end_critical(),
void  check_windows_running(void),
void  check_switcher_running(void),
#ifdef _cplusplus
}
#endif
#endif

```

Many readers of the first edition of this book noticed an error in the installation of TSR1XAMP.C. We were incorrectly calling `RestoreDosSwap` before having restored original interrupt handlers. Since `dos_setup` generates an INT 21h hook on 25h DPMI call we were messing up stacks; it potentially damages the state of DPMI after we have called `RestoreDosSwap`. We should call `RestoreDosSwap` only after all we are made at DPMI calls. This has been corrected in Listing 9-15.

Finally, there's `TSRUTIL.ASM` (Listing 9-17) which converts miscellaneous registers which we either didn't want to write in C, or couldn't.

### Listing 9-17: TSRUTIL.ASM

```

;TSRUTIL.ASM
;Define segment names used by C
;
_TEXT segment byte public 'CODE'
_TEXT ends
CONST segment word public 'CONST'
CONST ends
_BSS segment word public 'BSS'
_BSS ends
_DATA segment word public 'DATA'
_DATA ends
GROUP GROUP CONST, _BSS, _DATA
    assume CS:_TEXT, DS:GROUP
    public _new_int13, _inst_intr
IFDEF MULTI
    public _timer_int_chain
    public _new_int10, _new_int21, _new_int25, new_int26
ELSE
    public _deinstall
ENDIF
extrn  _ss_save near    ,save foreground SS
extrn  _sp_save near    ;save foreground SP
extrn  _unsafe_flag near ;if true, don't interrupt
extrn  old_int13 near

```

```

IFDEF MULTI
    extrn    _old_int8:near
    extrn    _old_int10:near
    extrn    _old_int21:near
    extrn    _old_int25:near    ; note difference between
    extrn    _old_int26:near    ; old_int25 and _old_int25*
    extrn    _dos_count:near
    extrn    _in_progress:near
ELSE
    extrn    _multiplex_id:near ;our INT 2fh id byte
ENDIF

_TEXT segment
IFDEF MULTI
,*****
;void far deinstall(void)
;function to use INT 2fh to ask TSR to deinstall itself
;the registers are probably all changed when our isr exits
;so we save them and perform the INT 2f. The TSR exit will
;eventually bring us back here. Then the registers are restored
;this function is called from the foreground, not the TSR
DEINSTALL    equ    1
_deinstall    proc    far
    push    si
    push    di
    push    bp
    mov     word ptr _ss_save,ss    ; save our stack frame
    mov     word ptr _sp_save,sp
    mov     cs,_ds_save,ds        ; save DS for later restore
    mov     bx,cs
    mov     dx,offset TerminateAddr ;bx:dx points to terminate address
    mov     ah,byte ptr _multiplex_id
    mov     al,DEINSTALL
    int     2fh                    ; call our TSR
;
;if ISR terminates ok, we'll skip this code and return to Terminate Addr
*
    jmp     short NoTerminate
TerminateAddr:
;Restore DS and stack
    mov     dx,cs _ds_save        ; bring back our data segment
    mov     ds,dx                ; destroyed by int 2f
    mov     si,2                  ; Set value for success
    mov     si,word ptr _ss_save   ; restore our stack
    mov     sp,word ptr _sp_save   ; destroyed by int 2f
NoTerminate:
                                ;Extend return value to word
    cbw
    pop     bp
    pop     di
    pop     si
    ret
_deinstall    endp
ENDIF
,*****
;void inc_unsafe_flag(void) - increment unsafe flag
,*****
inc_unsafe_flag proc    far
    push    ax
    push    ds
    mov     ax,DGROUP            ;make DS = to our TSR C data segment
    mov     ds,ax

```



```

inc     word ptr _unsafe_flag
pop     ds                    ;put DS back to whatever it was
pop     ax
ret

inc_unsafe_flag endp
;****
;void dec_unsafe_flag(void)  decrement unsafe flag
;****
dec_unsafe_flag proc  far
push   ax
push   ds
mov    ax,DGROUP             ;make DS = to our TSR C data segment
mov    ds,ax
dec    word ptr _unsafe_flag

pop    ds                    ;put DS back to whatever it was
pop    ax
ret

dec_unsafe_flag endp

;we can't trap the following interrupts in C for a number of
;reasons
; INT 13 returns info in the FLAGS, but a normal IRET
; restores the flags
;
; INT 25 & 26 leave the flags on the stack. The user
; must pop the off after performing an INT 25 or 26
;
; These interrupts pass information via registers such
; as DS. We don't want to change DS.
;
; Since DS is unknown, we must call the old interrupts
; via variables in the code segment. The _init_intr routine
; sets up these CS variables from ones with nearly-identical
; names in the C data segment in ISREXAMP.C.
;****
;void far init_intr(void)
;move interrupt pointer saved in the C program to our CS data area
;****
; note confusing distinction between e.g. _old_int13 and old_int13
SET_OLD MACRO _old_int, old_int
    tes bx, dword ptr _old_int
    mov word ptr cs:old_int, bx
    mov word ptr cs:old_int+2, es
ENDM

_init_intr proc far
push   es
push   bx

IFDEF M..T)
    SET_OLD _old_int10, old_int10
    SET_OLD _old_int21, old_int21
    SET_OLD _old_int25, old_int25
    SET_OLD _old_int26, old_int26
ENDIF
    SET_OLD _old_int13, old_int13
    pop    bx
    pop    es
    ret

_init_intr endp
;****
;void far new_int13(void) - disk interrupt
;****

```

```

_new_int13 proc far
    call    inc_unsafe_flag
    pushf
    call    cs:old_int13
    call    dec_unsafe_flag
    ret     2
;simulate interrupt call
; leave flags intact
_new_int13 endp

IFDEF MULTI
*****
;void far new_int21(void) - dos interrupt
*****
_new_int21 proc far
    sti
;see if function 0, if so jump to int 21h vector
    cmp     ah,0
    jne    int21_0
    jmp     cs:old_int21
int21_0:
    push   ds
    push   ax
    mov    ax, BGROUP
    mov    ds,ax
    cmp   word ptr [_in_progress], 0
    je    int21_1
    inc   word ptr _dos_count
;not in background, so skip next
;flag that the background has called dos
int21_1:
    pop    ax
    pop    ds
    pushf
    call   cs:old_int21
;simulate interrupt call
    pushf
    push  ds
    push  ax
    mov   ax, BGROUP
    mov   ds,ax
    cmp  word ptr [_in_progress], 0
    je   int21_2
    dec  word ptr _dos_count
;not r background, so skip next
int21_2:
    pop   ax
    pop   ds
    popf
    ret   2
_new_int21 endp
*****
;void far new_int10(void) video interrupt
*****
_new_int10 proc far
    call    inc_unsafe_flag
    pushf
    call    cs:old_int10
    call    dec_unsafe_flag
    iret
;simulate interrupt call
_new_int10 endp
*****
;void far new_int25(void) MS-DOS absolute sector read
*****

```

```

_new_int25 proc far
    call    inc_unsafe_flag
    call    cs:old_int25
    call    dec_unsafe_flag
    ret
; user must pop flags - MS-DOS convention
; so leave them on the stack
_new_int25 endp
;*****
;void far new_int26(void) - MS-DOS absolute sector write
;*****
_new_int26 proc far
    call    inc_unsafe_flag
    call    cs:old_int26
    call    dec_unsafe_flag
    ret
; user must pop flags - MS-DOS convention
; so leave them on the stack
_new_int26 endp
;*****
;void far timer_int_chain(void) - jump to next timer ISR
;we need to clean up the stack because of this call
;*****
_timer_int_chain proc far
    mov     _ax_save,ax
    pop     ax
    pop     ax
    mov     ax,_ax_save
    jmp     dword ptr _old_int8
_timer_int_chain endp
ENDIF

; save areas for original interrupt vectors
;
IFDEF MULTI
old_int10 dd 0 ;video
old_int21 dd 0 ;dos int
old_int25 dd 0 ;sector read
old_int26 dd 0 ;sector write
ENDIF
old_int13 dd 0 ;disk
_ds_save dw 0
_TEXT ends
_DATA segment
IFDEF MULTI
_ax_save dw 0
ENDIF
_DATA ends
end

```

### TSR Command Line Arguments

Any program built with the generic ISR can take optional command line arguments, but set its hkey and its multiplex inter-upt ID number. A command line option is also available to deinstall the ISR, the implementation of which will be discussed in detail later on.

[the command line syntax is: `tsrname [-k scan shift [-t multiplex id]] [-d deinstall]`. A valid shift status is any combination of

- 1 = Right Shift
- 2 = Left Shift
- 4 = CTRL
- X = All

Available multiplex ID is 00 through 1F, and specifies a unique INT 2F=ID starting at AH=00h.

The default hooker's MUX ID and the current multiplex ID's zero which turns into INT 2Fh Function 00h. Specifying a nonzero hooker's current multiplex ID's makes it possible for multiple run multiplex programs load simultaneously. INR File 16 demonstrates INR whose multiplex ID is not the default, you must specify the use of the other File of D's watches. For example:

```
C:\MUNDOC2\CHAPP9>tsrfile & 59 8
TSRFILE hooker is Alt-F1 (F1 decimal scancode is 59)
TSRFILE multiplex ID is default CDh

C:\MUNDOC2\CHAPP9>tsrrem & 60 4 F 1
TSRREM hooker is Ctrl-F2 (F2 decimal scancode is 60)
TSRREM multiplex ID is Cth

C:\MUNDOC2\CHAPP9>tsrrem -f 1 -d
TSRREM deinstalled
```

## Writing TSRs with the DOS Swappable Data Area

It is important to note that the DOS 3.0 through 3.3, as well as for DOS 5.0 and higher, and Function 5100h for DOS 4.x, give us access to the DOS SDA. This block of data contains some information which may be useful to the programmer. MS-DOS 1.0 SDA includes the current PSP, section of the free MS-DOS stacks, as discussed earlier, and as shown in gray detail in the Appendix. If the List of Lists List is using the DOS data structure SDA as DOS's data. For example, the value of the INT 2Fh Function 5100h (hex 02) returns the current PSP, where do you find the PSP's base from the SDA? The code example below illustrates DOS's three stacks by just hook a SDA's three pointers. I will present chapters as much extensive use of the SDA in a future implementation of my book. The SDA's large element of the DOS data segment \*In DOS 4.0 only, there can be multiple SDA's.

Although the actual need to write the SDA is less so for DOS 4.x only, the actual structure of the SDA is not a simple structure, so you must consult the computer's DOS 8.00 and 6.00.

What does the SDA offer to write INRs? If the File 16, it is the same time when they're doing get the current DOS's current INR's, we can see the current INR's use Functions 5100h and 5100h (hex 02) to get MS-DOS's stack and section of the SDA. It is always you to call MS-DOS at times with the definition of the current DOS's stack and section. We should not, that is, we do not need to see the SDA's structure, as it is not a simple structure. And, restoring the SDA can never be 100% reliable, since the SDA's base is not a simple structure, so you must consult the Appendix in the first edition of this book to encompass all of DOS's state.

There is one more thing to be aware of when writing INT 2Fh. When you're writing the INT 2Fh Function 80h, you need to be aware of the DOS's critical section. When DOS is in a critical section, you can not change the SDA's structure, so you must be aware of the use of the INT 2Fh Functions 81h or 82h. Note that the INT 2Fh Function 81h is MS-DOS's function, and application.

While it is possible to write a program that uses the INT 2Fh's, it is not a major drawback. It can be most effectively demonstrated by using a program that uses the function 5100h in conjunction with the INTRSPK program presented in Chapter 5.

### Listing 9 18: CRITSECT SCR Watches INT 2Ah Critical Section Calls

```
;; CRITSECT SCR -- Trap INT 2Ah critical section entry/exit calls
Intercept 2Ah
function 80h on_entry output "ENTER CRIT " at
function 81h on_entry output "EXIT CRIT " at
function 82h on_entry output "EXIT ALL CRIT
```

Running this script and leaving it loaded for as long as you like while running programs in native DOS will yield no results other than a lot of calls to INT 21h API 87h, which visits a critical section. This call comes from the INT 21h dispatch code we examined in chapter 6. But we never see any critical sections being entered! This is because, while DOS has code to generate these calls, at the appropriate times it is normally disabled, and requires patches to enable it. These patches, perhaps misnaming `gdi`, `crsdrv`, `cs`, `Windows` and `network` stubs perform the necessary patches and use the resultant calls in just the way we need to. In the case of Windows Iohack, it is likely that can be made visible inside a DOS box by adding the line `RollerDOSInt2A.Txt` to the `SYSTEM.INI` file. The CRITSEC.C script when then run in a DOS box will be much more than productive. The CRITSEC.C program in Listing 9-19 allows the necessary patches to be applied and removed by invoking the program with ON or OFF on the command line.

### Listing 9-19: CRITSEC.C

```
/*
CRITSEC.C
Andrew Schulman, January 1993
from Undocumented DOS, 2nd edition (Addison-Wesley, 1993)
Thanks to John Brennan (John.Brennan@uiowa.edu)
and to Norman D. Culver (CIS 70672,1257)
Reports on DOS critical sections, turns off or on
Table of critical section patch offsets at:
005:02C3h in DOS 3.1-3.3
005:0313h in DOS 4, 5, 6
to turn ON critical section calls in DOS
in DOS 3-4, poke in 50h (PUSH AX) to replace C3h (RET)
in DOS 3-6, poke in any non-zero value
*/
#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <dos.h>
typedef HK_FP
#define HK_FP(seg, ofs) ((void far *) (((unsigned long) (seg) < 16) | (ofs)))
#endif
// get DOS data segment
// cal. 21/52 (Get List of Lists), keep segment, discard offset
unsigned short get_dos_ds(void)
{
    _asm push es
    _asm xor es, es
    _asm mov es, ax
    _asm mov ax, 5200h
    _asm int 21h
    _asm mov ax, es
    _asm pop es
    // return value in AX
}
// get DOS version number
// if DOS 5+, use 21/3306 because of SETVER
void get_dos_version(unsigned char *pver, unsigned char *pmin)
{
    unsigned char verj, minj;
    // must cal 21/3306 before 21/30, but don't know if 21/3306 supported!
    // best bet is to clear BX beforehand, unsupported should set AL=FFh
    _asm xor bx, bx
    _asm mov ax, 3306h
}
```

```

_asm int 21h
_asm mov eax, b
_asm mov min, bh
if ((maj==0) && (min==0))
{
    _asm mov ax, 1000h
    _asm int 21h
    _asm mov ma, ah
    _asm mov min, ah
}
*pmaj = ma,
*min = min,
}

```

```
void fail(char *s) { puts(s), exit(1); }
```

```
typedef enum { STATUS, ON, OFF } REQU,
```

```
main(int argc, char *argv[])
```

```

{
    unsigned short far *patch_tab,
    unsigned short far *p,
    unsigned char far *fpop,
    unsigned short dos_ds,
    unsigned short tab_ofs,
    unsigned char dos_maj, dos_min;
    char *s;
    unsigned char op,
    REQU requ;

    // command line: ON, OFF, or (default) STATUS
    if (argc < 2)
        requ = STATUS;
    else
    {
        s = strdup(argv[1]);
        if (strcmp(s, "STATUS") == 0) requ = STATUS,
        else if (strcmp(s, "ON") == 0) requ = ON;
        else if (strcmp(s, "OFF") == 0) requ = OFF;
        else fail("usage: critsect [status | on | off]");
    }

    // get DOS data segment
    dos_ds = get_dos_ds();

    // get DOS version number program works with DOS 3.1-6.0
    // may have to check for DR-DOS, etc
    get_dos_ver(&dos_maj, &dos_min);
    if ((dos_maj == 3) | (dos_maj == 6) | (dos_maj == 5) && (dos_min == 0))
        fail("Unsupported DOS version");

    // get pointer to table
    tab_ofs = (dos_maj == 3) ? 0x02c3 : 0x0315,
    patch_tab = unsigned short far * 0x1000 + (dos_ds, tab_ofs),
    if (requ == STATUS)
        printf("Critical sectors patch table at %p\n", patch_tab),
#define RET 0xc3
#define PUSH_AX 0x50
    // walk table
    for (p=patch_tab; *p != 0; p++)
    {
        fpop = (unsigned char far *) RE_FP(dos_ds, *p);
        op = *fpop;
        if (dos_maj < 5)
        {
            if (op == RET) s = "OFF";
            else if (op == PUSH_AX) s = "ON",

```

```

else fail("Critical section patch table invalid"),
if (req == ON)      *fpop = PUSH_AF,
else if (req == OFF) *fpop = RET;
}
else
{
if (op == 0) s = "OFF";
else s = "ON",
// no way to check for valid/invalid,
// though anything other than 0 or 1 is suspicious

if (req == ON)      *fpop = 1;
else if (req == OFF) *fpop = 0;
}
if (req == STATUS)
printf("fpop = %02X %s\n", fpop, op, s),
}
if (req == ON) puts("Turned critical sections ON"),
else if (req == OFF) puts("Turned critical sections OFF"),
return 0;
}

```

Once DOS is generating INT 2Ah critical section calls, you will need to write an interrupt handler for INT 2Ah to keep track of critical sections. You can only swap the SDA when DOS isn't in a critical section.

Your INT 2Ah handler may be bypassed under Windows Enhanced mode because Windows patches DOS turning its INT 2Ah calls into far calls into the Virtual Machine Manager (VMM). Microsoft's *MS-DOS 5.0 Source: Reference* for MS-DOS 5.0 describes how to do this when it says that MRC 2 programs must acquire the Windows disk critical section by calling INT 2Ah AX=8000h and further more that calls to this function must be coded for a particular way because Windows expects this exact sequence" p. 446. The point is that Windows will patch its INT 2Ah AX=8000h so that it becomes a far call out VMM which implements its own Binary Critical Section and End Critical Section calls documented in the Windows DDK.

Returning to the SDA, INT 21h function 5100h in non-386 mode returns the following information:

```

05:51 points to 005 swappable data area
0X - size of area to swap when inDOS > 0
0K - size of area to always swap

```

INT 21h function 5100h for DOS 4 and higher returns a DS:SI pointer to an SDA list, which contains

Offset	Size	Description
00h	WORD	Count of SDAs
SDA ENTRY:		
02h	DWORD	Address of this SDA
06h	WORD	Data area length and type bit 15 - set if swap always clear if swap while inDOS > 0 bits 14-0 - length in bytes
08h		next SDA_ENTRY

To generate these functions give you pointers to the data areas that contains much of not all of the information related to the current process, as well as information specific to DOS, and that may be in progress. Since MS-DOS switches stacks when invoked with INT 21h, it is somewhat not recent. But since the stacks are part of this data area, you can save the current information

by moving the entire swappable data area and immediately call DOS without fear of trashing DOS's internal stacks. Unfortunately, there are also DOS variables that reside outside the part of the DOS data segment that is made available as the SDA.

In previous sections we checked the DOS flags to determine whether it was safe to activate and, once activated, we saved the current PSP, DDA, and extended error information. Using the DOSSWAP macro not eliminates these steps. You can determine whether it is safe to pop up by tracking the INT 2Mh critical section calls. If you are not in a critical section, it is safe to call DOS. If InDOS is active, just save the data area that is always swapped (up to 18 bytes). If InDOS is non-zero, then the swappable data areas must be saved (up to 1736 bytes, less than 2k).

The TSR saves the data to a memory block that was allocated during TSR initialization (see the InitDosSwap function in DOSSWAP.C). In our C TSR, this task must be performed before the SDA stack is allocated. Once the data area has been saved, we set the current PSP and DDA values to those for our TSR. When it is time for the TSR to exit, we just move back the SDAs that we've saved. Since this data block contains the current PSP, DDA, and extended error information, we don't need to deal directly with these values.

Note: A lot of this writing is not great, but that is the DOS SDA list returned by function \$1000h is static when MS-DOS is booted or whether it is changed dynamically during the course of MS-DOS execution. It appears that functions \$1000h and \$1001h are identical information for MS-DOS 4.x. Since function \$1000h was only for DOS 5, it shows a DOS 4.0 only phenomenon. I'd like to submit to totally guess if I could function \$1000h was used as another example of low IBM added over complexity to DOS when they got over heads over a revision 4.0.

The C module in Listing 9.20 contains functions for saving and restoring the DOS SDA.

### Listing 9.20: SDA Save and Restore Module DOSSWAP.C

```

/* DOSSWAP.C - Functions to manage DOS swap areas */

#include <stdio.h>
#include <dos.h>
#include <memory.h>
#include <tsr.h>
#include <put.h>

#define GET_DOSSWAP3 0x5d0e
#define GET_DOSSWAP4 0x5d0b
#define SWAP_LIST_LIMIT 20

struct swap_list /* format of DOS 4+ SDA list */
{
    void far* swap_ptr;
    int swap_size;
};

/* variables for 3.x swap work */
static char far * swap_ptr, /* pointer to dos swap area */
static char far * swap_save; /* pointer to our local save area */
static int swap_size_indos,
static int swap_size_4ways,
static int size;

/* variables for 4.x swap work */
static int swap_count; /* count of swappable areas */
static struct swap_list swap_list[SWAP_LIST_LIMIT];
/* list of swap areas */
static char far * swap_save[SWAP_LIST_LIMIT], /* our save area */
static int sup_flag[SWAP_LIST_LIMIT], /* flags if has been swapped */
static int dos_level, /* for level dependent code */
int dos_critical; /* in critical section, can't swap */
/*****

Function: InitDosSwap

```



Initialize pointers and sizes of DOS swap area.

Return zero if success

- 1 - no memory
- 2 - too many swap lists
- 3 - unsupported dos
- 4 - too many instance memory blocks

\*\*\*\*/

int InitDosSwap(void)

```
{
    union REGS regs;
    struct SREGS sregs;
    int ret;

    /* establish what dos level we're running
       make sure that DOSVER is not setup for this ISR
    */
    if ((_osmajor == 3) && (_osminor >= 10))
        dos_level = 3;
    else if (_osmajor == 4)
        dos_level = 4;
    else if (_osmajor == 5 || _osmajor == 6) /* 5 = 6 for now */
        dos_level = 5;
    else
        dos_level = 0;

    if (dos_level == 3 || dos_level == 5) /* use 215006 */
    {
        regs.x.ax = GET_BOSSWAPS,
        intdosx(&regs, &regs, &sregs);

        /* make sure no error occurred on interrupt function */
        if (regs.x.cflag)
            return(10);

        /* pointer to swap area is returned in DS:SI */
        FP_SEG(swap_ptr) = sregs.ds,
        FP_OFF(swap_ptr) = regs.x.si,

        swap_size_indos = regs.x.cx;
        swap_size_always = regs.x.dx,

        size = 0; /* initialize for later */
        swap_save = malloc(swap_size_indos);
        ret = ((swap_save == 0) ? 3 : 0);

        /* if we got mem, then see if we can add instance */
        if (!ret && add_instance_block(swap_save, swap_size_indos))
            ret = 4;

        return(ret);
    }
    else if (dos_level == 4) /* use Sd0b */
    {
        struct swap_list far *ptr,
        int far *iptr,
        int i;
        regs.x.ax = GET_BOSSWAP4,
        intdosx(&regs, &regs, &sregs);

        /* make sure no error occurred on interrupt function */
        if (regs.x.cflag)
            return(10);

        /* pointer to swap list is returned in DS:SI */
        FP_SEG(iptr) = sregs.ds;
        FP_OFF(iptr) = regs.x.si,
        swap_count = *iptr; /* get size of list */
        iptr++;
    }
}
```



```

        sup_list[i].swap_size),
    }
}
else
    return 1;
return 0;
}
/*****
Function: RestoreDosSwap
This function will restore a previously swapped dos
data area
*****/
void RestoreDosSwap(void)
{
    if (dos_level == 3 || dos_level == 5)
    {
        /* make sure its already saved and we have a good ptr */
        if (size && swap_ptr)
        {
            movedata(IP_SEG(swap_save), FP_OFF(swap_save),
                    FP_SEG(swap_ptr), FP_OFF(swap_ptr), size),
                size = 0;
        }
    }
    else if (dos_level == 4)
    {
        int i;
        for (i = 0; i < swap_count; i++)
        {
            movedata(IP_SEG(swap_save[i]), FP_OFF(swap_save[i]),
                    FP_SEG(swap_ptr), FP_OFF(swap_ptr),
                    sup_list[i].swap_size),
                sup_flag[i] + 0; /* clear flag */
        }
    }
}
}

```

For you to use the new DOSSWAP method for building TSRs (recompile TSR1X.AMP.C) Listing 9-14 with DDOSSWAP and then link with the DOSSWAP module. Listing 9-20. See the walkthru shown in Listing 9-1.

To use the DOSSWAP method in the multi-tasking, top pop up example presented later, we would need to save not only the foreground data area (data belonging to the process we are interrupting) but the background data area (our TSR's data) as well. The SDA for the TSR could be saved during TSR initialization. Using this method we would not need to deal with PSP, DDA and Extended Error values at all since they would already exist in the SDA. Always saving and restoring the data area may make it easier to design some sort of round robin task switcher. A hook would step through a number of independent applications. Interestingly, as we saw in a couple of the Microsoft Windows 3.x multitasker uses the undocumented SDA function. The "Origin of the SDA" discussion in chapter 8 points out that the SDA is not a genuine DOS integral data structure, but merely a portion of DOS's data signature that Microsoft exports for the benefit of networking software and other multitaskers.

If you ever examined the actual contents of the DOS SDA in our appendix, you can see that this area is quite large. Because of this you may need to weigh the advantages and disadvantages of using the SDA (the pros and cons) against using the SDA in TSRs is that you can actively do it at any time while DOS is busy (of course, a critical section has been flagged) using INT 21H. This would be most beneficial for a networking or round robin task switching, since the response to a task switch

and to be almost instantaneous. In our generic popup TSR, for example, using DOSSWAP allows us to popup instantly in the middle of a TYPE command.

There are two disadvantages. One is that Dos function captures memory to save the data areas. This has been avoided by leveraging the data to extended or expanded memory. Second, this technique is new, and its effectiveness has yet to be determined. It's something that you must play with and prove to yourself that it works.

### **TSRs and Task Managers**

A task manager is a resident program which is part of both Windows 3.11 (DOSMHE11) and provides control over a number of applications or tasks. Both Windows 3.11 and DOSMHE11 allow multiple sessions to be installed, giving the user control over some sessions' activity. When writing and using TSRs with a task manager, special considerations may be required, and we'll try to point out why.

First we'll go into a little more explanation of how a task manager works in general.

A task manager saves application program generally formatted session type of menu format. One menu item displays what applications can be started, while another menu may display applications that have already been started. It is from this list that the user can select which application to make active, with the task manager taking care of activation and deactivation of the chosen applications. When an application is deactivated, all information specific to that application is saved or swapped. A simple example would be that all of the application memory is copied and saved to a disk file or extended expanded memory. When it is time to activate an application, the saved information is reloaded in memory. In reality, the memory swapping is more subtle, just saving needed information such as program data and the MS-DOS swapped data area. Most times, the swapping is transparent to the application being activated or deactivated.

Okay, enough about task manager operations. What does this have to do with TSRs?

Most task managers allow multiple DOS COMMAND.COM sessions. Of course, only one session's running at any one moment, but you can still have four sessions to session.

TSRs are generally meant to popup once over the course of an application. However, the foreground application may run by one of many means, possibly each in its own DOS session. Because of this, there are two different means to work with the TSR. The first is that you could install the TSR in each task pending DOS session. Since each session preserves its own memory, each TSR instance would not interact with each other. But this can be expensive work for the TSR, since each TSR is attached to its own DOS session.

Now, what if you load the TSR into memory before starting the Task Manager Application. In this case, all DOS sessions have access to a single shared TSR. If the TSR is of the standard switching DOS session type, the TSR is active in the foreground session. Fortunately, the Windows and DOSMHE11 Task Manager provides means to install TSRs of more than one task creation and switching and also provides what is called instance memory.

Instance memory is a block of memory that is unique to each session. By keeping global data and the stack in this block of memory, each DOS session can have its own independent access to the TSR. The TSR provides a list of memory locations and addresses, and the Task Manager takes care of keeping the data unique for each session.

Creating a DOS session's own copy of the TSR data is not great if you want the TSR to provide a communications buffer or some other type of shared DOS sessions, but in general this is not how you would want the TSR to behave. Consider the information loaded before Windows. It has the variable CWD, the current working directory, which you know about declaring instance data. Then commands typed in any DOS session's window go where you hit the up arrow key in another DOS session, interesting behavior, certainly, but probably not what you want. The DOSKEY command line editor, in contrast, knows about instance data in books IN1-21h-AX-1605h and AX-4105h). As a result, even if DOSKEY is installed in Windows, commands typed in one DOS session do not leak across to the list in buffer screen in the DOS session.

The multiplex interrupt (INI 2fh) caused for communications to and from the task manager. Additional communication structures are also set up through an INI 2fh call. Since these functions are documented in MSDOS technical documentation, I will not go into a long-winded discussion on these items. The January, February, 1992 *Microsoft Systems Journal* has an excellent article by Douglas Holroyd on writing ISRs in an MSDOS 5.00 environment, which includes details on the task manager API. The Microsoft Developer Network (MSDN) CD-ROM also has an excellent article by David Long on "TSR Support in Microsoft Windows Version 3.1." TSR developers, even if they don't do Windows, must become familiar with these issues. See also your local computer dealer whether a standard user is going to load Windows after installing your TSR.

The following INI 2fh calls are offered by our TSR code. These calls are supported by Windows and DOSMIF11.

```
1605h  Windows Start
4801h  Build Notification Chain (attaches TSR to task manager)
4805h  Get Instance Data
```

When the task manager, such as DOSMIF11 or Windows 3.1 Standard mode, starts, it calls INI 2fh AX 4801h to build a chain of structures that indicates what programs are interested in receiving session information during session creation and swapping. The task manager also calls INI 2fh AX 4805h to get information about session instance data required by any programs. The 4805h call is very similar to the Windows INI 2fh AX 4805h initialization broadcast. In fact, DOSKEY uses the same piece of code to handle both 4805h and 1605h.

As our program processes the INI 2fh AX 4801h call, it provides the task manager with an address to a structure called SWCALLBACKINFO. This structure provides an address to code that processes session messages by the task manager. Thus, the task manager calls this address when it has information to pass on to the program. This allows the program to be aware of task status and events such as impending session swaps. If the session swap is a direct one, it is indicated by the notification message, also returns a value that prevents session swaps.

The source code module SWITCHER.C in Listing 9-21 manages part of the interface to the task manager. The structures were taken right from the MSDOS programmer's reference. None of this is undocumented, but to discuss ISRs without mentioning Windows task switching, and instance data, would be misleading.

### Listing 9-21: DOS Task Manager Interface Module SWITCHER.C

```
/* SWITCHER.C -- handles MS task switcher functions */
#include <stdlib.h>
#include <dos.h>
#include <memory.h>
#include <tsr.h>

/* The following structures are defined in the Microsoft "MS-DOS
Programmer's Reference" for DOS 5 */

/* SWINSTANCEITEM contains information about a block of instance data
You pass a list of these structure through the SWSTARTUPINFO
structure. The list of SWINSTANCEITEM structures is terminated
by a 32 bit (double word) 0 value. */
typedef struct {
    void far *pInstanceData, /* points to instance data */
    unsigned int iSize, /* instance data size */
} SWINSTANCEITEM;

/* SWSTARTUPINFO Contains information about a client program's instance
data. Used during an INT 2fh AX 4805h call */
typedef struct {
    unsigned int iVersion, /* ignored */
```

```

void far *sisNextDev, /* prev. handler's SWSTARTUPINFO */
void far *sisVirtDevFile, /* ignored */
void far *sisReferenceData; /* ignored */
SWINSTANCEITEM far *sisInstanceData; /* SWINSTANCEITEM structures */
} SWSTARTUPINFO,

/* SWAPIINFO contains information about the support level a client program
provides. Since our TSR does not provide any support, this block will
be set to zeros. Refer to Microsoft's technical reference for more
details on the values allowed in this structure. */
typedef struct {
unsigned int sisLength, /* size of this structure */
unsigned int sisAPI; /* API identifier */
unsigned int sisMajor, /* major version number */
unsigned int sisMinor, /* minor version number */
unsigned int sisSupport; /* level of support */
} SWAPIINFO,

/* SWCALLBACKINFO contains information about the client program */
typedef struct {
void far *scbiNext, /* next SWCALLBACKINFO structure */
void far *scbiEntryPoint, /* notification function handler */
void far *Reserved, /* reserved */
SWAPIINFO far *scbiAPI, /* list of swapping structures */
} SWCALLBACKINFO,

#define MAX_INSTANCE 20 /* limit to 20 instance blocks */

SWINSTANCEITEM SWInstanceItem[MAX_INSTANCE], /* instance blocks */
int currentInstance = 0, /* current number of blocks */

SWSTARTUPINFO SWStartupInfo,
SWAPIINFO SWSwapInfo,
SWCALLBACKINFO SWCallBackInfo;

int enhanced_windows = 1, /* indicating if we're running enhanced mode */
int switcher_critical = 0, /* indicating if a switch can happen */
extern union REGS regs,
extern struct SREGS sregs;

INJECT switcher_service,
extern notify_function(), /* used to get linkage into the supporting ASM
module. We just need the address */

/* the following is from NOTIFY.ASM */
extern void far process_existing_int2(INTRERRUPT REGS far *r, void far *next);
extern void far check_switcher(void far *ptr);

/* initialize Switcher Data structures.
This is where we configure the notification address and any instance
data that we might need. */
void init_switcher_structures(void)
{
/* zero out swap info structure since no API in this TSR support */
memset(&SWSwapInfo, 0, sizeof(SWSwapInfo));
/* set address of swap info in the call back structure */
SWCallBackInfo scbiAPI = &SWSwapInfo;
/* set address of notification function handler (in NOTIFY.ASM) */
SWCallBackInfo scbiEntryPoint = &notify_function;
/* zero out instance data structures for now, we'll add some later */
memset(&SWInstanceItem, 0, sizeof(SWINSTANCEITEM) * MAX_INSTANCE);
/* Set Structure ID */
SWStartupInfo.sisVersion = 3;

```

```

/* set up pointers to our instance data */
SWStartupInfo.sisInstanceData = SWInstanceItem,
)
/* add_instance_block() adds a block to the instance data array.
This must be called before isr goes resident
Returns 0 on success, or 1 if no more instance blocks available. */
int add_instance_block(void far *ptr, unsigned int size)
{
    int ret = 0;

    if (current_instance > MAX_INSTANCE-1)
        ret = 1;
    else
    {
        SWInstanceItem[current_instance].iisPtr = ptr;
        SWInstanceItem[current_instance].iisSize = size;
        current_instance++;
    }
    return(ret),
)
/* process_switcher_int2f() checks to see if a switcher command has
been passed. If yes, it will process it and return a 0 indicating
to the caller not to continue with int 2fh processing. A return
of 1 indicates the caller should process the 2F interrupt. */
int process_switcher_int2f(INTERRUPT_REGS far *r)
{
    int ret_code,
    segread(&sregs),

    switch (r->eax) /* see if we're interested */
    {
        case 0x1605 /* windows initialization */
            enhanced_windows = !(r->edx & 1), /* zero in enhanced windows */
            /* call other 2f handlers */
            process_existing_int2f(r, &SWStartupInfo.sisNextDev),
            /* set up saved registers with required address */
            r->es = sregs.ds;
            r->bx = (unsigned int) &SWStartupInfo,
            ret_code = 0; /* tell caller that all is processed */
            break,

        case 0x4b01: /* build notification chain */
            /* call other 2f handlers */
            process_existing_int2f(r, &SWCallBackInfo.scbiNext);
            /* set up saved registers with required address */
            r->es = sregs.ds,
            r->bx = (unsigned int) &SWCallBackInfo,
            ret_code = 0, /* tell caller that all is processed */
            break,

        case 0x4b05: /* get instance data */
            /* call other 2f handlers */
            process_existing_int2f(r, &SWStartupInfo.sisNextDev),
            /* set up saved registers with required address */
            r->es = sregs.ds;
            r->bx = (unsigned int) &SWStartupInfo;
            ret_code = 0, /* tell caller that all is processed */
            break;

        default /* let every thing else fail through */

```

```

        ret_code = 1,
        break;
    }
    return(ret_code);
}
/* Signal Windows to not task swap */
void windows_begin_critical()
{
    if (enhanced_windows)
    {
        regs.x.ax = 0x1681,
        int86(0x2f, &regs, &regs);
    }
}
/* Signal Windows that task swap ok */
void windows_end_critical()
{
    if (enhanced_windows)
    {
        regs.x.ax = 0x1682,
        int86(0x2f, &regs, &regs);
    }
}
/* See if windows is running in enhanced mode.
   If not, decrement enhanced mode flag */
void check_windows_running(void)
{
    regs.x.ax = 0x1600, /* windows check */
    int86(0x2f, &regs, &regs);
    if (regs.h.al == 0) /* windows not installed */
        enhanced_windows = 0;
}
/* See if task switcher running (done in assembly) */
void check_switcher_running(void)
{
    check_switcher(ESMCallBackInfo);
}

```

As stated previously, the task manager sends messages to our program through a notification call. These parameters are passed through registers and, because notification is not through an interrupt, the INTR pins will not be called through interrupt handling functions. That is, because you can't control the register values and registers containing parameters may be used by C before the program gets a chance to process the call, NOTIFY.ASM and using 9.32 manages the rest of the interface to the task manager. It also uses the I/O port routine called process\_running and 9.31 in this file. Process\_running and 9.31 passes the INTR messages to other programs that may be interested in it. The program does not use a critical section since we register information after the call to the other INTR 21h handlers return. The only documented API usage for the task manager is that the program calls which programs exist and their addresses of the task manager INTR 21h function.

### Listing 9-22: DOS Task Manager Interface Module NOTIFY.ASM

```

,NOTIFY.ASM
/
/ Routines to handle notification messages from the task switcher
/ these are not easily (if possible at all) written in C since the
/ parameters are passed in registers, and NOT called through an
/ interrupt (if they were interrupt functions and a C interrupt function
/ could handle them. Most of these functions are stubs, returning
/ all is ok. The only thing we might want to prevent is switching
/ during a TSR critical section, say if we were handling a time

```



```

; sensitive piece of hardware.
;
; None of the notify functions are called from the C program, and should
; not be'
;
; Define segment names used by C
;
_TEXT segment byte public 'CODE'
_TEXT ends
CONST segment word public 'CONST'
CONST ends
_BSS segment word public 'BSS'
_BSS ends
_DATA segment word public 'DATA'
_DATA ends
BGROUP GROUP CONST, BSS, _DATA
    assume CS:_TEXT, DS:BGROUP
    public _notify_function
    public _process_exiting_int2f
    public _check_switcher

    extrn _switcher_service:near ,switcher service function
    extrn _switcher_critical:near ,critical = 1, means no swap'
    extrn _old_int2f:near ,original INT 2fh vector

_TEXT segment
;****
;function called by task switcher to notify us of switcher states
;****
_notify_function proc far
    push    ds
    push    bx
    push    ax
    mov     ax, BGROUP ;setup to our C data segment
    mov     ds,ax
    pop     ax

;
; save current switcher service function. Documentation says this
; may change, so save every time you get one
;
    mov     word ptr [_switcher_service], di
    mov     word ptr [_switcher_service+2], es

;
; setup to call notification function through a jump table defined
; below in data segment
;

    cm     pas,7 ;make sure valid function value
    ja     notify_exit

    mov     bx,ax
    shl     bx,1
    call   [SwitcherTable+bx]

notify_exit:

    pop     bx
    pop     ds
    ret

```

```
_notify_function endp
```

```

;*****
; RET_ZERO used for:
; function 00 - Initialize Switcher
; function 03 - Activate Session
; function 04 - Session Active
; function 05 - Create Session
; function 06 - Destroy Session
; function 07 - Switcher Exiting
;*****
RET_ZERO proc near
    xor     ax,ax             ;signal ok to process
    ret
RET_ZERO endp

```

```

;*****
; Function 01 - Query if OK to suspend session
;*****

```

```

QuerySuspend proc near
    mov     ax, word ptr _switcher_critical
    ;; Ooops! According to MS-DOS Programmer's Reference for
    ;; DOS 6.0 Task Switcher API Patch, pp 495-496, under
    ;; v 5 of the DOS task switcher, the QuerySuspend handler
    ;; must install a patch that preserves CX.
    ret
QuerySuspend endp

```

```

;*****
; Function 02 - Suspend Session
;*****

```

```

SuspendSession proc near
    mov     ax, word ptr _switcher_critical
    ret
SuspendSession endp

```

```

; Call Old INT 2fh function used to process switcher INT 2fh functions
; from our C code. Its main reason for being is to put
; the INT 2fh registers back to the way when we received the interrupt.
; This is so that the function we call has them in the right place

```

```

; void process_existing_int2f(INTEERRUPT_REGS far *r, void far *next)

```

```

;*****
_process_existing_int2f proc far

```

```

    push   bp
    mov    bp,sp

```

```

; Move the original INT 2fh address into the code segment
; we'll lose the data segment when we reset the registers

```

```

    mov    ax, word ptr _old_int2f
    mov    word ptr cs:[_save_int2f], ax
    mov    ax, word ptr _old_int2f+2
    mov    word ptr cs:[_save_int2f+2], ax

```

```

; save address of SSTART.PINFO.next

```

```

    lds   bx, [bp + 10]    ; get address we want to put ES (cx):BX into
    mov   cs:[tmp_date1], bx ; save address of STRUCTURE
    mov   cs:[tmp_date2], ds

```

```

; get stack back to just after all registers were pushed by the C
; interrupt routine

```

```

;
lds    bx, [bp+6]    ;get pointer to regs structure
mov    ax, ds
cli
mov    ss, ax
mov    sp, bx
sti

; Restore regs as they were in original INT 2fh call
;
IFDEF TURBOC
    pop    bp
    pop    di
    pop    si
    pop    ds
    pop    es
    pop    dx
    pop    cx
    pop    bx
    pop    ax
ELSE
    pop    es
    pop    ds
    pop    di
    pop    si
    pop    bp
    pop    bx
    pop    bx
    pop    dx
    pop    cx
    pop    ax
ENDIF

ENDIF

; save INT 2fh function for later compare
;
mov    cs:tmp_data3, ax ;save request

; let others process request
;
push    dword ptr cs:_save_int2f

; now update the structure with the address returned in ES:BX
;
push    ds
push    ax
push    bx
mov    ax, 0GROUP    ;get data seg so we can store ES:BX
mov    ds, ax
pop    ax ;save value of bx from INT 2fh call

mov    bx, cs:tmp_data1 ;get address of structure
mov    [bx], ax ;store ES:BX values
mov    [bx + 2], es

;
; if function 4b01, then structure address is same as next address
;
cmp    cs:[tmp_data3], 4b01h
je     skip_dec
dec    bx ;point to start of structure

```

```

    dec     bx
skip_dec
    push   ds
    pop    es

    pop    ax
    pop    ds
    iret

_process_exiting_int2f endp

; ****
; check if task switcher already running.
; if so, setup call back structure to hook notification chain
; ****
_check_switcher proc far
    push   bp
    mov    bp,sp
    mov    bx,0           ;documentation says must be zeros
    mov    di,0
    mov    es,di
    mov    ax, 4b02h
    int    2fh
    or     al,al
    jnz   no_switcher ;none installed, so just exit

; save switcher service function
;
    mov    word ptr [_switcher_service], di
    mov    word ptr [_switcher_service+2], es
    les   di, [bp+4] ;get address of callback structure
    call  dword ptr [_switcher_service]
no_switcher:
    pop    bp
    ret
_check_switcher endp

_save_int2f dd 0
tmp_data1 dw 0
tmp_data2 dw 0
tmp_data3 dw 0
_TEXT ends

_DATA segment
;
; jump table used to dispatch switcher notification messages
;
SwitcherTable dw offset RET_ZERO      , InitSwitcher (00)
               dw offset QuerySuspend , (01)
               dw offset SuspendSession , (02)
               dw offset RET_ZERO      , ActivateSession (03)
               dw offset RET_ZERO      , SessionActive (04)
               dw offset RET_ZERO      , CreateSession (05)
               dw offset RET_ZERO      , DestroySession (06)
               dw offset RET_ZERO      , SwitcherExit (07)
_DATA ends
end

```

## Removing a TSR

Before we get started, we would like to extend special thanks to Tom Paterson for all the help he provided in this section for the first edition of *Undocumented DOS*.

There are times when you need to remove a TSR from memory—TSRs use memory, and at times conflict with other software. Usually, the TSR can be removed so that there is no trace of it as if it had never been installed; the removal process unlinks the TSR from the interrupt vector table and returns all TSR memory back to MS-DOS.

One situation that can arise when trying to remove a TSR is that other programs can chain into the same interrupts that the TSR is processing. This prevents you from removing the TSR from the interrupt chain. Consider the following TSR example that chains into interrupt 8:

In this process, the code saves the current interrupt address—the `INT8` save data storage—to be called from inside the TSR. The interrupt vector table now points to the TSR. The following diagram describes the chaining process before and after a TSR adds itself into the chain:

```

Before our TSR loads
INT 08 -----> Original Interrupt Service Routine (ISR)
After our TSR is resident:
INT 08 -----> TSR ISR -----> Original ISR
  
```

At this point, if you wanted to remove the TSR from the interrupt chain, you would just put the original ISR address, saved earlier, back into the interrupt vector table. The TSR is out of the link. But what if another program has placed itself into the interrupt chain, as in the following diagram?

```

INT 08 -----> NEW ISR -----> TSR ISR -----> Original ISR
  
```

If you put the value stored for the original ISR back into the interrupt vector table, you would be cutting off the new program that has inserted itself into the chain. When it is your time to remove your TSR from the interrupt chain, therefore, check to make sure the interrupt table is pointing to the correct `INT8` interrupt service routine. If it is, then you know that no other TSR has been loaded after your ISR. If it isn't, what then?

Correctly, there are two new strategies that have been discussed for use at installation. If a TSR that has interrupt handlers is the middle of a chain, neither of them can do itself as a TSR programming in C:

1. If you can't remove a TSR because there is another application in the chain, then it just removes itself the 2 possible. Basically, create all your original ISR that just jumps to the original ISR. On TSR removal, de-allocate all memory except for the section the serial ISR still uses. This retains the interrupt chain integrity.
2. IBM has proposed a standardized format for interrupt handlers. All handlers will start with a very strict set of followed by data that contains a signature and the original ISR address. The purpose is to get over the serial data block and continue with interrupt processing. This works well with assembly language, since you have control over how your ISR appears. This does not fit easily into C, though, since a programmer has no control over the structure of the C interrupt function. To achieve this standard in a high-level language, all interrupts would have to be generated to code in an assembly module. The C programmer would have to store the original interrupt addresses into the CMM segment of the assembly module. The assembly module would jump to the associated C interrupt blocks. The proposed interrupt structure is shown in Figure 9-23.

**Figure 9 23: Proposed Interrupt Chain Structure**

```

new_interrupt:      jmp _interrupt_handler
previoushandler:   dd 0
signature:         dw 4240h ; *KB*
hardwareflag:     db 0
                  jmp hardwarereset
                  db 7 dup (0) ; Reserved

```

The signature field is an *aid* for other programs in identifying this interrupt structure. The hardware flag field is important for interrupt handlers that will be issuing a I/O hardware operation. This field should be zero unless the ISR handles a hardware interrupt and is the FIRST installed ISR which becomes the last in the chain if other handlers are installed. The *jmp* instruction to a hardware reset is mainly not needed for software interrupts.

The following code snippet demonstrates how this would be done. Don't forget that the ISR initialization section would need to save the real interrupt address in the `int8` variable contained in the assembly language routine.

```

public _new_int8, _old_int8
extern _process_int8
_new_int8 proc far
    jmp far _process_int8 ;go to C handler
_old_int8
signature
hardwareflag
    jmp _hardwarereset ;if hardware function
    db 7 dup (0) ; Reserved
_hardwarereset:
    retf ;just return in case someone calls into this label
_old_int8
_new_int8 endp

```

The interrupt handler in C would look like it always does.

```

void interrupt far process_int8(INTERRUPT_REGS Z)
{
    // interrupt processing
    _chain_intr(_old_int8)
}

```

Note, this is only a proposed standard, our sample ISR does not include such code, but it probably should.

The generic ISR developed in this chapter can be triggered by a command line argument. If you use `process` loaded ISRs in FAT, the command `ISRI11 D` loads the resident copy of `ISRI11.FAT`, the `D` tells the generic ISR shell to start a resident ISR program to unload and not to load up another copy.

The generic ISR uses `INT8h` to communicate between the newly executed copy and the resident copy. The address function number `MFxval` can be set with a command line parameter, but the substitute is `MF` in handovered to the ISR program. A value of zero in the `MF` register is an install check and a value of one is a kernel command. The use of zero for the install check is dictated by the standard Multiplex Interrupt interface.

The following steps are required for installation. These steps follow the `INT8h` deinstallation call. Step one must be performed by the ISR, since only it knows what addresses to restore to the interrupt vector table. Steps two through five can be performed either by the ISR or the calling application.

- 1 See if all of the interrupt vectors the TSR has changed still point to the interrupt service routines. Restore all vectors that have not changed since the TSR changed into them. If any interrupt vectors have changed, disable the TSR operations and skip the following steps to leave the TSR in memory.
- 2 Set the parent PSP of the TSR to be that of the application that executed the INT 2Fh deinstallation call. The parent PSP's value is contained at offset 10h of the TSR's PSP.
- 3 To return control to the calling application, set the termination address to be a location in the calling application.
- 4 Set the current PSP to the PSP of the TSR. Doing this allows DOS to call the TSR's memory and close any open files.
- 5 Execute the normal DOS terminate function INT 2Fh function 4Ch. When the TSR has been terminated by DOS, control returns to the terminate address set in step three.
- 6 At this point, all registers have unknown values. The calling process must save them before issuing the INT 2Fh deinstallation command. They can now be restored.

If the TSR is not performing steps two through five, it must return its PSP value as part of the INT 2Fh interface.

The generic TSR performs steps one through five. INT 2Fh ASM contains the deinstallation call through INT 2Fh, and INT 2Fh AMP contains the rest of the code that handles deinstallation. The pseudocode in Figure 9-24 describes the generic TSR operation.

**Figure 9-24: Generic TSR Unload Pseudocode**

Foreground Application	Resident Application (TSR)
Save SI, DI, BP, SS and SP registers	
Set AX=001h	
Set BX:DS to terminate address (TERM_ADDR, used in step 3 above)	
Perform INT 2Fh	TSR receives INT 2Fh request
	Set up local stack
	Attempt to unlink interrupts
	If unable to unlink all interrupt vectors, then:
	Set AL = 0uff
	Return from INT 2Fh
If we get here, the deinstallation failed	
Check AL register and display proper message	
EXIT	Get PSP (foreground app)
	Set parent PSP field in our PSP
	Set terminate address in our PSP to address passed via INT 2Fh
	Set PSP to TSR's PSP
	Call MS-DOS terminate
TERM_ADDR	
Restore SI, DI, BP, SS and SP registers	
EXIT	

Many times a TSR is not able to deinstall due to interrupt vectors changed by another application. A return from the INT 2Fh call with AL = 0ffh indicates that the TSR could not unlink itself from one or more interrupts. Deinstallation also fails if the TSR is not installed to begin with. This is indicated by AL = 00h on return from the INT 2Fh call.

## Sample TSR Programs

To create the generic TSR, we built three different simple TSRs, two of which are pop-up versions of programs from other parts of this book.

### TSRFILE

The first simple program, `TSRFILE` (see Listing 9-25), demonstrates that you really can make DOS hit Ctrl-C a 50th or 99th activation calls while you're popped up in the middle of some other program. When `TSRFILE` pops up, it prompts the user for a filename and then displays the file on the screen. No screen saving/restoring amenities are provided.

`FILE.C` uses the Microsoft C++ dos functions for file I/O and memory allocation. These functions translate directly into the appropriate INT 21h calls and are thus preferable to using the intdos or intdosfar macros. Instructions for converting `FILE.C` into `TSRFILE.EXE` are found in the maketds shown in Listing 9-1. It can also be compiled as a .com and given `FILE.COM`.

### Listing 9-25: Example Pop-up FILE.C

```

/* FILE.C */
#include <dos.h>
#include <conio.h>
#include <fcntl.h>
#include <share.h>

char file_prompt[] = "file ";
char cant_open[] = "Can't open file!\n";
char error_reading[] = "Error reading file!\n";
char insuff_mem[] = "Insufficient memory, Press any key...\n";
char ctrl[] = "\n";

#define PUTSTR(s) \
    dos_write(STDERR, (char far *) s, sizeof(s)-1, &count)

#define MIN_PARAS 4
#define WANT_PARAS 64
#define BYTES (paras << 4)
#define STDERR 2

#define TSR
void application(void)
#else
main(void)
#endif
{
    char buf[81];
    char far *s;
    unsigned rcount, wcount, ret, paras, seg,
        int i;

    /* prompt for filename */
    if (PUTSTR(file_prompt) != 0)
        return;

    /* get filename */
    if (( dos_read(STDERR, buf, 80, &rcount) != 0 ) & (rcount < 5))
        return;

    /* replace CRLF with NULL */
    buf[rcount-2] = '\0';

    /* try to allocate: first try a lot, then a little */
    if (_dos_allocmem(WANT_PARAS, &seg) == 0)
        paras = WANT_PARAS;
    else if ( _dos_allocmem(MIN_PARAS, &seg) == 0)
        paras = MIN_PARAS;
    else

```



```

{
    PUTSTR(Insuff_mem);
    return;
}
FP_SEG(s) = seg;
FP_OFF(s) = 0;
/* open file */
if (_dos_open(buf, O_RDWR | SH_DENYNO, &f) != 0)
    return PUTSTR(cant_open);

/* display file */
while ((ret = _dos_read(f, s, BYTES, &rcount)) != 0) {
    if (_dos_write(STDERR, s, rcount, &ccount) != 0)
        break;
    /* write one more CRLF */
    PUTSTR(crlf);
    if (ret)
        PUTSTR(error_reading);
    /* free memory */
    _dos_freemem(seg);
    /* close file */
    _dos_close(f);
    PUTSTR("Press any key...");
}

```

The C makes two stabs at allocating memory, because it pops up over COMMAND.COM there is almost no free memory available. COMMAND uses the largest block in memory, and all that are left are little dubs and dubs like the environment we freed during TSR initialization. You can see this situation when running our second example, the MEM program as a TSR.

### TSRMEM

One problem with the MEM program presented in Chapter 7 was that since it was a stand-alone program, you could only examine the memory map from within MEM itself. By putting the generic TSR and MEM together to form TSRMEM.EXE, you can examine the memory map with other programs. For example, you can clearly see how COMMAND grabs the largest chunk of memory, leaving almost nothing free.

```

C:\NUNDO2\CHAP9>tsrfile
C:\NUNDO2\CHAP9>tsrmem -k 59 8 -f 1
C:\NUNDO2\CHAP9>\sidekick\sk

[Hit TSRMEM hotkey, Alt-F1]
seg  Owner      Size
09F3  0008  00F4 ( 3904)  config [15 48 67 ]
0AE8  0AE9  00D3 ( 3376)  0BC1 c:\dos33\command.com [22 28 ]
0B8C  0000  0003 (   48)  free
0BC0  0AE9  0019 (  400)
0B0A  171A  0006 (  208)
0B8E  0000  0000 (   0)  free
0BE9  0BEA  0575 ( 22352)  EP1 FA ]
115F  1160  0589 ( 23440)  -k 59 8 -f 1 [1B 23 24 2F f4 F5 ]
1719  171A  1905 (105296)  0BDB C:\SIDEKICK\SK.COM [0B 09 10 13
                               16 1C 21 25 26 2B ]
30CF  0AE9  8730 (553728)  [30 F8 ]

```

The largest block in the MCB chain, totaling 87,308 paragraphs, is not marked "free." Instead, it is owned by PSP0AE9, looking back along the MCB chain, which also functions as a PSP chain, you can see that 0AE9 is COMMAND.COM. In fact, there are only three paragraphs of free memory located directly after COMMAND. Even the environment, TSRMEM freed was possessed by Side-

back for use as its environment. If you hit `INSRIII`'s `break` `Ah D` at this point, it will ask for the filename and then report "Insufficient memory."

However, if you hit the `COMMAND` by running an application and then hitting `Ah D`, there is no kernel panic of any kind because `COMMAND` is no longer hogging the largest block. Again, this shows a problem for you hit `INSRIII`'s `break` within some other application. The MCB being on top of `COMMAND.COM` has now been replaced by the following:

```
300F 300E 0000 ( 208)
300D 300E 2920 (168656) 3000 c:\apps\EPSILON.EXE (00 05 16 )
3400 0000 50F4 (384832) free (30 F8 )
```

Now, thanks to `INSRIII`'s paragraphs of free memory, you now have no trouble allocating memory in `INSRIII`.

Putting `MEMM` to use the generic `ISR` was straightforward. We had to get rid of a call to `calloc`, replace a few calls to `exit` with simple returns, expand any " " characters into "0x", and take a few other minor adjustments. The key change, however, was to mix with a version of `printf` that does not call `_scanf` after initialization. The `ISR` would have no near heap from which to allocate memory.

This non-`malloc` version of `printf` is provided in the module `PUT.H`, which can be used with any program that uses the generic `ISR`. The non-`malloc` version of `printf` uses the `stdarg` facilities of `VSNI.C` in support of the `va` function pointer, which can help create functions that take variable arguments. `PUT.H` also contains a collection of other helpful functions. Prototypes for the functions appear naturally in `PUT.H` (see Listing 9-26).

#### Listing 9-26: No-`malloc` `STDERR` Output `PUT.H`

```
/* PUT.H -- STDERR output routines, no malloc */

// calls _dos_write, returns number of bytes actually written
unsigned doswrite(int handle, char far *s, unsigned len);

// displays ASCII string on STDERR
unsigned put_str(char far *s);

// displays character on STDERR
unsigned put_char(int c);

// displays number (width, radix) on STDERR
unsigned put_num(unsigned long u, unsigned wid, unsigned radix);

// PUT.H includes alternate version of printf goes to STDERR,
// if desired use #ifdef. Same prototype as <stdio.h>

// get string from STDERR, returns actual length
unsigned get_str(char far *s, unsigned len);

#define putstr(s)      ( put_str(s); put_str("\n"); )
#define put_hex(u)    put_num(u, 4, 16)
#define put_long(u)   put_num(u, 9, 10)
```

In `MEMM`, the `if (perror == 0) { ... }` statement was used to conditionally compute either a standard or a `ISR` version. For example:

```
#ifdef TSR
void fail(char *s) { printf("Zs\n\n", s); return; }
#else
void fail(char *s) { puts(s); exit(1); }
#endif
```

The changes needed to make MEM a pop-up were all of a similar nature and are so straightforward and uninteresting that we leave them as an exercise for the reader. In any case, the existing TSRMEM.EXE can be found on the disk that accompanies this book.

## TSR2E

Finally, we jump the gun a little bit by porting a program from the next chapter in this book. As Jim Kule will explain, INT\_21H is the backdoor to the DOS command interpreter. The INT2E command interpreter from Chapter 10 can be easily turned into a TSR to provide a custom pop-up copy of COMMAND.COM.

This really does work. The only peculiarity is that on occasion when you pop up TSR2E and type in a command, the following message appears from the resident portion of COMMAND.COM:

```
Memory allocation error
Cannot start COMMAND, exiting
```

Note, however, that this is different from the horrifying message one sees when the MCB chain has been trashed:

```
Memory allocation error
Cannot load COMMAND, system halted
```

The message "exiting" rather than "system halted" is for real. If you try to execute the command again it works. In any case, it should probably come first of other INT\_21H services found in the next chapter.

You, of course, TSR2E out of your assembly text or batch commands such as DIR or COPY, but also of course external commands (not in other programs or even batch files).

As shown in the booklets presented earlier in the chapter you build TSR2E by concatenating the generic TSR components with Jim Kule's source files INT21.C, HWI2E.ASM, and IO21.ASM. These files are unchanged. All changes for a simple TSR are confined to the module INT21.C, where we change the name of the module's export from main to application, add a test that sets TSR2E as a pop-up of which COMMAND.COM is already running, and use the facilities of the PC's code loader (not the C standard library's malloc and free facilities). The alternate version of INT21.C appears in Listing 9-27.

### Listing 9-27: Example Pop-up TEST2E.C

```
#include <stdlib.h>
#include <string.h>
#include <dos.h>
#include "put.h"

#define WK_FP(seg,ofs) \
    ((void far *)(((unsigned long)(seg) < 16) ? (ofs)))

extern unsigned foreground_psp, // in TSRXAMP.C
extern int Send2E(char *command), // in SEND2E.C

static char buf(80);
static int running = 0;

typedef enum { SAVE=0, RESTORE } SAVEREST;
typedef void (*INTVECT far *INTVECT)();
void Interruptal(int restore)
{
    static INTVECT int_7b, int_23, int_24;
    if (restore)
    {
        _dos_setvect(0x7b, int_7b);
        _dos_setvect(0x23, int_23);
        _dos_setvect(0x24, int_24);
    }
}
```

```

)
else
{
    int_1b = _dos_getvect(0x1b);
    int_23 = _dos_getvect(0x23);
    int_24 = _dos_getvect(0x24);
}
}

void application(void)
{
    // don't run if we are already running
    if (running)
        return;
    running++;

    // don't execute INT 2Eh if COMMAND.COM already running
    // see if COMMAND.COM running by checking if current PSP is the
    // same as its own parent (see chapter 10)
    if (foreground_psp ==
        *((unsigned far *) MK_FP(foreground_psp, 0x16)))
    {
        put_str("COMMAND.COM already running");
        running--;
        return;
    }

    put_str("FSR COMMAND SHELL. type DOS commands, or BYE to quit\r\n");
    for (;;)
    {
        put_str("$ ");
        if (! get_str(buf, 80))
            break;
        if (strcmp(buf, "bye") == 0 || strcmp(buf, "BYE") == 0)
            break;
        interrupts($SAVE);
        Send2F(buf);
        interrupts($RESTORE);
    }
    putstr("Bye ");
    running--;
}
}

```

We save and restore the Ctrl C, Ctrl Break, and Critical Error interrupts around the call to Send2F(). With this precaution, even Ctrl C, Ctrl Break, and Critical Errors are handled properly within the INT 2Eh pop up.

```
FSR COMMAND SHELL type DOS commands, or BYE to quit
$ dir a
```

```
Not ready error reading drive A
Abort, Retry, Fail? a
```

```
$ dir *.* /w
```

```
Volume in drive C is BARANJIAN
Directory of C:\UNDOC\RMICHELS
```

```
DOSSWAP C      EXTERR C      ^C
```

```
$ bye
Bye
```

One word about trying to install a TSR from within the pop up command interpreter: DON'T! This will hang your system sometime after you exit the pop up.

## Multitasking TSR

Finally, let us discuss TSRs that don't pop up at a user hotkey, but which do their work in the background. We call such programs multitasking TSRs to distinguish them from pop-ups. MULTIC is a multitasking TSR that can be modified to perform multiple background tasks, from disk file copying to background communications.

The example presented here is an enhancement to the DOS PRINT utility. It periodically activates by INT 8 and INT 20h searches a SPPOOL directory for files having the extension SPL. When it finds a match, the TSR uses PRINT's INT 21h function 01h interface to ask ERUN1 to print the file. Once the file has been submitted for printing, the TSR periodically obtains a status report from PRINT. If the file is no longer in the print queue, its printing is complete, the file is deleted.

```
C:\VND0C2\CHAP9> print
C:\VND0C2\CHAP9> mults
C:\VND0C2\CHAP9> copy \vndoc\rmichels\* .asm \spool\* .spl
C:\VND0C2\CHAP9> dir \spool
```

```
Volume in drive C is RAMANUJAN
Directory of C:\SPOOL
.                <DIR>          3-23-89    9 54p
..               <DIR>          3-23-89    9 54p
TSRUTIL.SPL     5852    9-17-90   11 40p
STACK.SPL       1758    9-17-90   11:36p
4 File(s)      96208 bytes free
```

```
C:\VND0C2\CHAP9> print
C:\SPOOL\TSRUTIL.SPL is currently being printed
```

```
C:\VND0C2\CHAP9> dir \spool
Volume in drive C is RAMANUJAN
Directory of C:\SPOOL
.                <DIR>          3-23-89    9 54p
..               <DIR>          3-23-89    9 54p
STACK.SPL       1758    9-17-90   11 36p
3 File(s)      102400 bytes free
```

Basically, the program manages its independent processes, the foreground process and the background process. When a hotkey is hit or a process is started, the current process is suspended. Its registers are saved in the stack for later restart. Once the TSR has been asked or restarted, the current instance of the suspended process's current context is restored and the program execution continues.

Thus multitasking is achieved by maintaining a context-based on the time slice of CPU. Each process gets a specific amount of time, the preceding example, the foreground process gets the most time slice to take care of performance. It is possible to add more tasks, but one would need to maintain a list of SS:SP sets to service all running processes.

Most of the code is similar to INT8XAMP. The major differences are that what was thought to be an interrupt hotkey and that instead of completing its work during activation, the TSR is suspended for later restart. Because DOS is not recursive, the TSR still must off by the rule of not interrupting DOS when it is active.

In most computer systems, multitasking is a method of quickly switching from one task to another so that the computer appears to be running multiple tasks at the same time. Of course, it is not as simple as that. Multitasking systems are designed so that resources such as disks, screens, and keyboards can be shared by multiple applications. True multitasking systems such as OS/2 supply interface routines that are recursive. The code segment of each routine is its own instance, and this code segment can be started in multiple processes at the same time. Each user can run one or will have a new instance of DOS. But as we know, the INT 21h API is not recursive so multitasking example is controlled to some extent by the DOS flags.

## Task Switching

Any process has what can be called its context or frame. This consists of the following items:

- Register values, including code, data, and stack segment values
- Program Segment Prefix (PSP) or process ID
- Disk Transfer Address
- Extended Error Information

To stop a task or context switch, these items must be saved and replaced by ones that pertain to the new task. The example simplicity drops between two tasks. If more independent tasks were required, we would find it necessary to store information for each task. Each item in the list might contain the relevant BIOS information (INT 13h, etc.) and a pointer to the process stack segment and offset. Your high-level task manager should switch to a copy of the SDA for each task. This code to perform the list management is more complex and is not presented here.

Perhaps the most striking presentation is called time slicing because each task gets a predetermined amount of time when to run. If needed, more intelligence could be added that would control the percentage of time each process gets. This could be based on usage of the operating system (INT 11h), or by task (INT 13h). You could fabricate these interrupts, and if a process is making extensive use of the resources, you could lower its time slice to give other processes more time.

This method is used during the background process. When the file being printed is still in the print queue, a timer (timer variables) set so that background processing terminates immediately. There is no point in running in the background if the program is simply waiting for the print spooler to complete its job.

One word regarding the use of MULTI with other task switchers: DOS 1+ things can get very confusing with multitasking on top of multitaskers. MULTI is not compatible with Microsoft Windows. However, Windows/Multi intercepts the Windows startup message (INT 21h AX=1605h) with the help of a special value in the CX register. This tells Windows to abort its startup.

## MULTI Installation

MULTI uses some file techniques already described. Some interrupt handling routines must be installed, so setup is necessary because the central interrupt process. On returning from most interrupt calls, the files are restored to their condition just before the interrupt was invoked. But three of the interrupts we are interested in handling (INT 13h, 20h, and 26h) do not follow this convention. INT 13h, like INT 13h, creates a new condition with flags, so an INT 13h handler must end with a RET instruction. Like INT 13h, INT 13h, while the DOS bootstrap task routines (INT 25h and 26h), leave the EAX register intact, so they exit with a RET rather than an IRET. This code is marked as #DDEMULTI (see the end of the ASM source code for more info).

Another difference between TSR/VAMP and MULTI is that in MULTI, the address of the timer interrupt is stored in the stack that is created in the main procedure. This is typical of interrupt routines, so the timer interrupt will occur at this address when the background process is first activated.

## Timer Interrupt

MULTI is related to the timer interrupt, of course, but by the INT 8 timer tick interrupt. MULTI's INT 8 handler increments a counter, counting count, which keeps track of how many timer ticks have occurred. It compares that value to a value called a given number of ticks. Once the current tick count exceeds the process limit, a count of 0 is written, except within INT 28h, in which it will be one, but must not be more than one. The timer routine (pages 145-146) that process suspends and the other process activates.

At this point, you can suspend the process. Depending on what process you are currently executing, you can call either suspend background or suspend foreground. To suspend the foreground, set the

stack to be the local ISR stack. In either case, carry out the save/restore register used earlier in the function as INTRXAMP, setting INI 1Bh, 33h, and 34h, and swapping PSP, DVA, and extended error information.

One factor to get to the new INTR function is restores the stack to one half page above the foreground process and passes control along the INI 8 into the program.

One special condition applies the first time you activate the background process. At this point, you have never interrupted this process, so its registers are not on the stack. Because of this, an assembly language function timer out is an assembly language jump to the next timer interrupt handler.

We rather referred to a flag variable called `isactiveflag`. Since there is no INI 80h flag to use, you have to create your own. This flag is set TRUE when critical BIOS services, INI 10h, disk, and INI 13h disk, or the BIOS bootstrap disk services, INI 35h and 36h, are in progress.

### Idle Interrupt

The timer interrupt is not the only way you can interrupt the background. You also use the BIOS idle interrupt, INI 28h, `int28h` service. This allows you to continue processing while the system is waiting at the COMMAND prompt. Otherwise, you would never return to the background while COMMAND was awaiting orders.

During an INI 28h `int28h` BIOS interrupt function returns FALSE, and the background is not already active, you set the foreground `int28` to zero and the `int28` active variable to TRUE. Then wait for the `int28` active variable to go FALSE before continuing. This allows you to interrupt the task switch to the background process.

### Keyboard Interrupt

Notice that we have installed a service routine for INI 9, the keyboard hardware interrupt. When ever the user presses a key, the background task `int9` is set to zero, causing the task to go into a suspended state more quickly. This gives the user better response time.

### Printing

The main loop of a function is our example is the stack-stand process. This function performs the work of searching for files and submitting a function to the system using INI 21h function 011. Notice that the example calculates a reverse loop that decreases during the main loop function. This is to avoid constant disk access with the background task. The background `int9` is also set to zero in a number of places to ensure that the background becomes suspended at that time.

### MULTI.C

Listing 9-28 shows MULTI.C, the source code for the multitasking. SR-10 and MULTI.FBI use the instructions found in the `makefile` in Listing 9-1.

#### Listing 9-28: MULTI.C

```

/* MULTI.C */
#include <stdio.h>
#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <dos.h>
#include <io.h>
#include "tsp.h"
#include "put.h"

#define SEARCH_DIR    "C:\\$POOL\\*"
#define STACK_SIZE    4096 /* must be 16 byte boundary */
#define SET_DTA        0x1a /* SET Disk Transfer Address */
#define GET_DTA        0x2f /* GET Disk Transfer Address */

```

```

#define BACKGROUND_TICKS 2
#define FOREGROUND_TICKS 16
#define BACKGROUND_YIELD 0
#define FOREGROUND_YIELD 0

struct prReq
{
    char level;
    char far *fname,
},

char far *stack_ptr; /* stack for our background TSR */
char far *ptr,
unsigned sb_save, /* slot for stack segment register */
unsigned sp_save, /* slot for stack pointer register */
unsigned unsafe_flag = 0; /* set true by various interrupts */
    nt first_time ^ 1, /* flag for first time in running background */

int my_psp; /* our TSR's psp */
int foreground_psp, /* PSP of interrupted foreground process */
int foreground_dta_seg, /* DTA of interrupted foreground process */
int foreground_dta_off,
int ctr=0,
int tic_count = 0; /* counts timer ticks */
int in_progress = 0, /* true if we're in background process */
int breakState, /* Extended Break Status of foreground */
    nt dos_count 0, /* count of our calls into DOS */
int multi_critical = 0; /* indicates we can't swap right now */
char search_work[65],
struct ExtErr my_errinfo;
struct ExtErr foreground_errinfo;

int foreground_limit = FOREGROUND_TICKS, /* foreground cycle limit */
int background_limit = BACKGROUND_TICKS, /* background cycle limit */

char search_dir[65] = (SEARCH_DIR), /* dir to search for spool files */
volatile int int_2B_active 0, /* true if activated by int 2Bh */
volatile int interval_timer, /* for sleeping a number of ticks */

/* old interrupt pointers are stored here */
INTVECT old_int8, old_int9, old_int10, old_int13,
INTVECT old_int1B, old_int21, old_int23, old_int24, old_int25,
INTVECT old_int26, old_int28, old_int2f,

/* prototypes for this module */
void main_loop();
void interrupt far new_int8(INTERRUPT_REGS),
void interrupt far new_int9(INTERRUPT_REGS),
void interrupt far new_int10(void),
void interrupt far new_int13(void),
void interrupt far new_int1B(INTERRUPT_REGS),
void interrupt far new_int21(INTERRUPT_REGS),
void interrupt far new_int23(INTERRUPT_REGS),
void interrupt far new_int24(INTERRUPT_REGS),
void interrupt far new_int25(void),
void interrupt far new_int26(void),
void interrupt far new_int28(INTERRUPT_REGS),
void interrupt far new_int2f(INTERRUPT_REGS),
int spooler_active(void);
int search_spl_quel(char * fname);
void suspend_foreground(void);
void suspend_background(void);

/* returns nonzero if PRINT installed */
int spooler_active()
{
    union REGS regs;
    regs.x.ax = 0x0100; /* PRINT install check */
    intB6(0x2f,&regs,&regs); /* call multiplex interrupt */
}

```



```

return(regs.h.al == Doff); /* FF if installed */
}
/* returns nonzero if file is in the spooler queue */
int search_spl_queue(char * fname)
{
    union REGS regs,
    struct SREGS sregs;
    char far * que_ptr,
    char que_name[65];
    int i;
    int found = 0;
    if (spooler_active())
    {
        regs.x.ax = 0x0104, /* get spooler status */
        int86x(0x2f,&regs,&regs,&sregs);
        /* on return from call DS:SI points to print queue */
        FP_SEG(que_ptr) = sregs.ds,
        FP_OFF(que_ptr) = regs.x.si,
        /* release hold on spooler, side effect of status */
        regs.x.ax = 0x0105,
        int86x(0x2f,&regs,&regs,&sregs),
        while (*que_ptr && !found) /* while stems in queue */
        {
            for (i = 0; i < 64; i++)
                que_name[i] = *(que_ptr + i);
            if (found = !strcmp(que_name,fname))
                break;
            que_ptr += 64;
        }
    }
    return(found);
}

void main_loop()
{
    struct find_t c_file,
    union REGS regs;
    struct SREGS sregs,
    struct prReq prRequest;
    struct prReq far * ptr;
    int sleep_cntr;
    while (1)
    {
        strcpy(search_work,search_dir),
        strcat(search_work, " SPL"), /* create dir search string */
        interval_timer = 78 * 30, /* search every 30 seconds */
        while (interval_timer) /* wait between each dir search */
            background_timer = BACKGROUND_TIMER, /* yield for fgnd */
            if (!_dos findfirst(search_work,_A_NORMAL,&c_file))
            {
                /* if spooler installed, dos 3.xx+ and file size > 0 */
                if (spooler_active() && _osmajor >= 3 && c_file.size)
                {
                    strcpy(search_work,search_dir),
                    strcat(search_work,c_file.name); /* full pathname */
                    prRequest.level = 0,
                    prRequest.fname = search_work;
                    regs.x.ax = 0x0101,
                    ptr = &prRequest,
                    sregs.ds = FP_SEG(ptr);
                    regs.x.dx = FP_OFF(ptr);
                    int86x(0x2f,&regs,&regs,&sregs),

```



```

/* put back original PSP */
SetPSP(foreground_psp);

/* put back original INTS */
_dos_setvect(0x1b,old_int1b);
_dos_setvect(0x23,old_int23);
_dos_setvect(0x24,old_int24);

/* get error info */
GetExtErr(&my_err)info;
SetExtErr(&foreground_err)info;

tic_count = 0;
in_progress = 0;
int_28_active = 0;

foreground_limit = FOREGROUND_TICKS; /* set default limit */

/* put back extended break checking the way it was */
SetBreak(breakState);
)
/*****
 * TIMER TICK INTERRUPT HANDLER
 *****/
void interrupt far new_int8(INTERRUPT_REGS r)
{
    _enable(); /* enable interrupts */
    tic_count++;

    if (interval_timer)
        interval_timer--;

    if ((in_progress &&
        !multi_critical &&
        dos_count == 0 &&
        ((tic_count >= background_limit &&
         !DosBusy()) && !unsafe_flag) ||
        (int_28_active && !Int28DosBusy()) &&
         tic_count >= background_limit))
    {
        suspend_background();
        restore_stack();
    }
    else if (!(in_progress &&
        ((tic_count >= foreground_limit &&
         !DosBusy()) && !unsafe_flag) ||
        (int_28_active && !Int28DosBusy()) &&
         tic_count >= foreground_limit))
    {
        set_stack();
        suspend_foreground();
        if (first_time)
        {
            first_time = 0;
            timer.int chain();
        }
    }
    old_int8(); /* call old handler */
}
/*****
 * KEYBOARD INTERRUPT HANDLER
 *****/
void interrupt far new_int9(INTERRUPT_REGS r)
{
    unsafe_flag++;
    old_int9();
}

```

```

if (!in_progress)
    background_limit = BACKGROUND_YIELD; /* set to swap to bgnd */
foreground_limit = 18, /* since user hit keyboard */
unsafe_flag--;
}
/*****
* CTRL-BREAK INTERRUPT HANDLER
*****/
void interrupt far neu_int1B(INTERRUPT_REGS r) { /* do nothing */ }

/*****
* CTRL-C INTERRUPT HANDLER
*****/
void interrupt far neu_int23(INTERRUPT_REGS r) { /* do nothing */ }

/*****
* CRITICAL ERROR INTERRUPT HANDLER
*****/
void interrupt far neu_int24 (INTERRUPT_REGS r)
{
    if (_osmajor >= 3)
        r.ax = 3; /* fail dos function */
    else
        r.ax = 0,
}

/*****
* DOS IDLE INTERRUPT HANDLER
*****/
void interrupt far neu_int28(INTERRUPT_REGS r)
{
    if (!in_progress && !int28dosBusy() && !unsafe_flag &&
        tic_count > foreground_limit)
    {
        foreground_limit = FOREGROUND_YIELD; /* stop foreground */
        int_28_active + 1;
        _enable(), /* STI */
        while (!int_28_active)
            ; /*spin waiting for task swap to bgnd*/
    }
    (*old_int28)(); /* call old handler */
}

/*****
* DOS MULTIPLEX INTERRUPT HANDLER
*****/
void interrupt far neu_int2f(INTERRUPT_REGS r)
{
    /* See if windows is starting up. If so, then set CX to be non-zero
    to keep windows from starting. This TSR and Windows are not
    compatible. */
    if (r.ax == 0x1605) /* windows init msg */
        r.cx = 1;
    else
        chain_intr(oid_int2f), /* pass int to other handlers */
}

main()
{
    unsigned memtop, dummy,
    void far* far* tmpptr,
    puts("Multi-Tasking PRINT spooler installing ),
    if (_osmajor < 3)
    {
        puts( Error: MS-DOS version 3.00 or greater required");
    }
}

```

```

    exit(1);
}

if (! spooler_active)
    puts("Warning: Print Spooler not active");

InitInDos(),
my_psp = GetPSP();

/* MALLOC is stack for our TSR section */
stack_ptr = malloc(STACK_SIZE);
stack_ptr += STACK_SIZE;
ptr = stack_ptr;
*(--stack_ptr) = 0xf2, /* set up stack as if an IRE was done*/
*(--stack_ptr) = 0xd2;
stack_ptr -= 4;
tmp_ptr = stack_ptr;
*(tmp_ptr) = main_loop;

/* get interrupt vectors */
old_int8 = _dos_getvect(0x08), /* timer int */
old_int9 = _dos_getvect(0x09), /* keyboard int */
old_int10 = _dos_getvect(0x10), /* video int */
old_int13 = _dos_getvect(0x13), /* disk int */
old_int21 = _dos_getvect(0x21), /* dos int */
old_int25 = _dos_getvect(0x25), /* sector read int */
old_int26 = _dos_getvect(0x26), /* sector write int */
old_int28 = _dos_getvect(0x28), /* dos idle int */
old_int2f = _dos_getvect(0x2f), /* dos multiplex int */

init_intvect(), /* init asm variables */
_dos_setvect(0x08,new_int8),
_dos_setvect(0x09,new_int9),
_dos_setvect(0x10,new_int10),
_dos_setvect(0x13,new_int13),
_dos_setvect(0x21,new_int21),
_dos_setvect(0x25,new_int25),
_dos_setvect(0x26,new_int26),
_dos_setvect(0x28,new_int28),
_dos_setvect(0x2f,new_int2f),

#define PARAGRAPHS(x) ((FP_OFF(x) + 15) >> 4)

/* release unused heap to MS-DOS */
/* A: MALLOCs for TSR section must be done in TSR_INIT() */
/* calculate top of memory, shrink block, and go TSR */
segread(&regs),
newtop = sregs.ds + PARAGRAPHS(ptr) - _psp,
_dos_seth.ack(newtop, psp, &dummy),
_dos_keep(0, newtop),
}

```

When compiling the ALL THE TSR example under Microsoft C 6.00 we take a two-pass compilation "go.go". Notice that the new int28 function simply sets up the tick counting variables so that the next word becomes active. The code then sets the int\_28\_active semaphore and waits for it to be cleared.

```

int_28_active = 1;
while (int_28_active)
;

```

The idea behind my code is to wait until a task swap occurs. During the task swap int\_28\_active is set to zero.

But when compiling with full optimization, the compiler sees that int\_28\_active is initialized to 1 and is not altered in the while loop. Therefore it figures this is an infinite loop and generates a warning.

What it does not know is that the timer interrupt routine will clear this flag. To avoid this problem but still allow optimization, declare the `int_28_active` variable with a "volatile" attribute:

```
volatile int int_28_active;
```

The `VOLATILE` keyword tells the compiler that the variable may change from an external source.

This multitasking TSR is a simple example to which a few enhancements could be added. Memory could be swapped to disk or EMS as needed. The TSR could provide a pop-up service where background operations can be controlled dynamically. If keyboard and CUI/O is required by the background task, you must be careful to save and restore the appropriate settings. You also must not switch tasks while in the middle of a BIOS video service. For a true multitasking system, MS-DOS alone is probably not the way to go. Commercial multitaskers such as Windows or DesqView give you much more functionality. But such systems use many of the same principles outlined here.

## Command Interpreters

by Jim Kyle

Every operating system that permits more than a single program to run requires some sort of command interpreter. In the earliest days of computing, the human operator interpreted commands by pecking out the correct program orders and deck and setting the system in action. Now a program—often called the shell because it surrounds the system kernel, not more formally termed the command interpreter—does the job. The command interpreter prompts the user for input and then reacts to that input.

For DOS, exactly a command interpreter is the closest contact they ever have with an actual operating system. The familiar prompt from COMMAND.COM (the character-based command-line shell that comes with MS-DOS) is almost universally called the DOS prompt. Although the interpreter is now in fact part of MS-DOS itself, alternate interpreters such as 4DOS.COM and the MKS Korn Shell are also available.

This chapter examines first the basic requirements that must be met by any command interpreter. The standard COMMAND.COM (on DOS 5.0) provides specific examples of some requirements. While examining these requirements, the chapter explores some related and somewhat surprising that DOS provides to simplify the task of interfacing another command interpreter. The first section concludes with some shell programs (see *Appendix A*) to replace COMMAND.COM (or, program files) to see exactly what the requirements are for creating a command interpreter.

When we discuss the requirements established in this chapter for any COMMAND.COM successor, it will discuss these requirements. Every section will begin above the environment, how it works, and how to use COMMAND.COM, and will work with the user's shell section on the desktop. In addition, it will discuss the various routines for locating and dealing with environment blocks, since the environment strings contain essential information used by the standard interpreters, such as the prompt and the path to the interpreter itself.

Next we examine some alternative command interpreters that are now available together with some shells that are virtually self-sufficient to COMMAND.COM.

The chapter concludes with some programs that combine the use of downlevel and uplevel environment features to permit editing of either the target environment or the currently active copy of it. This program, ENVEDIT, works with any command interpreter that supports the `set` and `setenv` DOS (and INTRIN) features. The command interpreter need not provide full support, so long as it merely returns enough buffer for the service.

Several other DOS-related activities aspects of the DOS programmer's interface are covered in this chapter: using the back door to the command interpreter (INTRIN) and the DOS master environment block (the `env` and `kernel` that quite important internal command interface—NINTRIN Function A'h) is also discussed in detail.

## Inside COMMAND.COM

The major requirements for any command interpreter are as follows:

- To provide a means of obtaining commands from the human operator
- To act on them by dispatching appropriate processes

An additional requirement is that these actions be carried out in a loop so that more commands may be issued after all currently issued commands have been processed.

Before we look at the details of these requirements, let's see just what happens inside COMMAND.COM itself. We'll put these details into perspective. The summary that follows resulted from an NRSMP script using DOS 5.0. The script reported a total of INI 21h and INI 2Eh. I then edited the NRSMP report to include only one cycle around COMMAND.COM's inner loop and to remove excessive detail that obscured the nature of the actions. I also added comments to the report at the top of each line to show what COMMAND.COM is doing at each point.

The cycle begins a new return to COMMAND.COM's loop following execution of the previous command:

```
ret from child process
INT 2F at FFFF D55E, AX=4000 Get exit code Exit Code 00h, Exit Type 00
INT 2F at FFFF 94D5, AX=4800 ALLOC ffffh paras FAIL (0008), only 89EDh availab.e
INT 2F at FFFF 94EE, AX=48AD ALLOC 89EDh paras returned seg 161fh
INT 2F at FFFF:B725, AX=2522: Set INT22 ==> 0B40 D17F
INT 2F at FFFF:B72A, AX=2523: Set INT25 = 04A0 D14B
INT 2F at FFFF:B731, AX=2524: Set INT24 ==> 0B40 D156
INT 2F at FFFF B00A, AX=3E24 Close 0005 through 0015
```

COMMAND.COM first saves the exit code returned by the previous process and stores it in an internal variable. COMMAND.COM then scans back all available RAM that the previous process took over. It does this by using the DOS 16-bit terminate function. This is apparently done so that COMMAND.COM will own the next available location if it needs to reload that part of itself. Next, the three interrupt vectors for the INT are restored to point to routines within COMMAND.COM. Next, a set of 16 bit variables, one for each of the 16 interrupt vectors, are stored.

Note that one of the previous calls to INI 21h came from the HIMM; this could indicate that they're related to the DOS 16-bit terminate function. However, a similar trace (not shown), run from 4DOS.COM, showed that COMMAND.COM reported a different sequence of calls to do the same thing, so it's clear that this sequence is not part of DOS itself. As you shall see later, COMMAND.COM automatically puts portions of its code in the HIMM if you run DOS-HIGH and that's where these calls originate.

With cleanup from the previous process nearly done, COMMAND.COM next resets its internal error flags with one exception: it does not reset the INI 2Eh AX=1234E. This is followed by calls that appear to be support routines that reset up facilities and finally by output of a CR LF pair to re-claim the ST 300 handle. The command line words are then part of the trace identify the line generated at ST 300; the interrupt routine is also called on exit:

```
INT 2F << 9672 5119, AX=122E DOS HOOK GET/SET ERROR TABLE
== 2F ==> AX=122E, BX=0014, CX=0000, DX=5100
INT 2F << 9672-5128, AX=122E DOS HOOK GET/SET ERROR TABLE
== 2F ==> AX=122E, BX=0014, CX=0000, DX=5E02
INT 2F << 9672 5137, AX=122E DOS HOOK GET/SET ERROR TABLE
== 2F ==> AX=122E, BX=0014, CX=0000, DX=5E04
INT 2F << 9672-5146, AX=122E DOS HOOK GET/SET ERROR TABLE
== 2F ==> AX=122E, BX=0014, CX=0000, DX=5E06
INT 2F << 9672 5175, AX=122E DOS HOOK GET/SET ERROR TABLE
== 2F ==> AX=122E, BX=0014, CX=007B, DX=5E08
```



```

121 at 9672:5183, AX=6300 Get Lead Byte Table* 9672_5E08
121 at 9672:0199, AX 3800 Get/Set Country Code: 9672_985A
121 at 9672_5418, AX=4000 Write to 0001_00h 0ah

```

The execution is complete, so `COMMAND.COM` generates a prompt for display, building it from the current drive code and current working directory to satisfy my "Sp5g" PROMPT specification. This is written to `STDOUT` and, as `COMMAND.COM` the call to `check_cwd` (see Chapter 8) since a documented function `INT 21h AX=FD08h and 8A9Fh` and in case a redir needed just to be able to be flushed and restarted. I removed call lines from the following excerpt:

```

121 at 9672_01EE, AX=1900 Current drive is_ 03
121 at 9672_1FD3, AX=475C Get CWD for 00 returned JB052\1SPY
121 at 9672_5418, AX=4000 Write to 0001_0 8 JB052\1SPY
121 at 9672:3E48, AX=023E: Display Char '*'
121 at 9672_0240, AX 5009 Flush redir printer output
INT 2F << F0CB_450C, AX=1125 REDIRECTOR Get/Set stream state
121 at 9672_0282, AX=5008 Set redir printer-01
INT 2F << F0CB_450C, AX=1125 REDIRECTOR Get/Set stream state

```

With the prompt displayed `COMMAND.COM` then calls `DOSKEY` using `INT 21h` to find the keyboard buffer. If `DOSKEY` is active. Since `DOSKEY` is not active, this call has no effect and the next call to `21/0A` is necessary to get the typed input. With `DOSKEY` active, the call to `21/0A` would be skipped over. `COMMAND.COM` spends most of its execution time only to either `DOSKEY` or `21/0A`, waiting for keystrokes.

```

INT 2F << 9672_028C, AX=4810 DOSKEY Get input
== 2F >> AX 4810, BX 858C, CX 00F9, DX=8B59
121 at 9672_02C4, AX=0A10: Buffered input to 9672_8B59

```

When I press `ENTER` to end the prompt control returns to `COMMAND.COM`. The program extracts a `CR` if part of a acknowledge, but then parses the command line to determine what to do next.

```

121 at 9672_5418, AX=4000: Write to 0001_00h 0ah
121 at 9672_0313, AX=2901 Parse FCB (mode 01) from 9672_8BDE to 9672_8E05
121 at 9672_039C, AX=2901 Parse FCB (mode 01) from 9672_0081 to 9672_005C
121 at 9672_0381, AX=2901 Parse FCB (mode 01) from 9672_0081 to 9672_006C

```

The interaction between terms is what I DID. The string `1SPY` in this example, represents the name of a single program installed, compared to the discussion later in this chapter. `COMMAND.COM` makes two calls to `21/AE00` (see below) to parse out other command interpreters use only a single call.

```

INT 2F << 9672_2BFC, AX=AEO0 INST CWD Check of 0101
CH=FF
== 2F >> AX AE00, BX 8BDC, CX FF00, DX=FFFF, returns 00
INT 2F << 9672_2BFC, AX=AEO0 INST CWD Check of 0101
CH=00
== 2F >> AX AE00, BX 8BDC, CX 0000, DX FFFF, returns 00

```

The return value of 00 from both these calls tells the command interpreter that `DHDI` is not an installed command. `COMMAND.COM` then examines its stored internal commands and does not find `DHDI`, then exits. The next step is to try to locate a file of that name in the current working directory. The computer call to `21/0B` by Redirector routine is built into the `21/0B` to return `APPEND` as a new reference to a file on drive happens to be made.

```

121 at 9672_32ED, AX 1A00 Set OTA to 9672_96D5
121 at 9672_34E9, AX=4700 Get CWD for 00 returned JB052\1SPY
121 at 9672:34F4, AX=4E00 FindFirstz didit.???

```

```
INT 2F << F0CB AB45, AX:1123 REDIRECTOR Qualify name: d:\dit ???
name not resolved returned 0000
```

SYSTEM INTERRUPT 0000 indicates that such a file exists. With the return value non-zero, COMMAND.COM would repeat the previous step on each directory named by the PATH environment variable. However, in this case, DIBBLE exists in the current working directory (CWD), so COMMAND.COM searches for DIBBLE.COM, DIBBLE.EXE, or DIBBLE.BAT. Comparison of the file type bits with the DIBBLE file attributes indicates that the file found is DIBBLE.COM, so no additional search succeeded. COMMAND.COM releases the resources it grabbed earlier than early 3.14p0's root test, so DIBBLE.COM as a child process. Again, an immediate call to the redirector appears:

```
INT 2F at 9672:5562, AX:4700 Get (MD for 04 returned UDOS2\ISPY
INT 2F at 9672:29FF, AX:490B FREE seg 161FH
INT 0F 0B4D 01DC, AX:4B00 Loading D:\UDOS2\ISPY\DIBIT.COM
INT 2F << F0CB AB45, AX:1123 REDIRECTOR Qualify name D:\UDOS2\ISPY\DIBIT.COM
name not resolved
```

DIBBLE.COM now performs only one action: printing the words "Loaded and run", then it lets the system INTR interrupt terminate the normal DOS termination function. Since our script does not suspend the interrupt, the action of termination comes from the interrupt to the DOS redirector for the usual file cleanup and closing of remote files. Then the program returns from the INTR interrupt, suspending our execution until the next loop and ending our example:

```
INT 2F at 162B:07, AX:0900 Print string Loaded and run
INT 2F << F0CB 93E2, AX:1127 REDIRECTOR Cleanup at termination
INT 2F << F0CB 6DC5, AX:131B REDIRECTOR Close all remote files
set from child process
```

At this point the system returns control to us as we started. The shell retrieves the exit code, grabs memory of a remote process, cleans up, and restores the interrupt vectors, although the DIBBLE.COM program may change from COMMAND.COM as well as all other command interpreter processes, always the principle of loading the same program as always, even if it never did them in the first place.

Don't let the redirector give you a false sense of private networking capabilities; see Chapter 8 for a complete description of its abilities and where. Later in this chapter we look at command-line redirection.

To summarize, the command interpreter shows a prompt, gets operator input, determines how to respond to that input, and does so by either running an installed command, running an internal command, or spawning a child process, or a combination of these. When one of those three actions is selected, the interpreter enters the loop; the interpreter retrieves the exit code for an external process, cleans up, restores the interrupt vectors, and does it over again. The key MS-DOS functions involve the 0Ah system file-based input and 4Bh read and execute process.

## Requirements of a Command Interpreter

### Obtaining Operator Input

The most essential requirement for any command interpreter is that it obtain commands to be interpreted, without that nothing else has any meaning.

You may see that operator input can be obtained from keystrokes entered directly by the user in response to a prompt from the command interpreter. There are, however, several alternate methods, all of them frequently used.

**The DOS Prompt** The command interpreter usually signals the user that it's waiting for input by issuing a prompt message. Usually this is simply the "C prompt." Some more sophisticated programs turn on the cursor or the mouse pointer to indicate that input is needed, but—unlike other such programs—DOS displays that information when seeking input, so that the mere presence of the cursor on the screen serves as the prompt.

In an environment that multitasks, such as Windows, the prompt is always present because input can be accepted even when other programs are running. In this case the mouse pointer often changes to an hourglass shape when input cannot be accepted due to the temporary existence of a system modal condition.

In the more conventional command-line operation, the prompt consists of a relatively short sequence of characters. The default prompt message of COMMAND.COM is simply C, where C indicates the drive letter of the current drive and is simply a usual denotation. Virtually all organized interpreters, though, give the user a means to modify the prompt into whatever form might be desired.

Although a dedicated PROMPT command seemed to define custom prompt messages, the actual message is stored as a string of characters in memory space and can be changed in the same manner as any other environment string. Many third-party users use PROMPT to specify other than the default C prompt and to accomplish other tasks. One example of COMMAND.COM logic showed how this prompt is produced. In addition, ANSYS can be used to put the current drive and directory at the top left corner of the screen; for example, when this screen prompt appears in its normal page:

**Keystrokes** The normal source of input to the command interpreter is the system keyboard, but it is available only as a last resort. The interpreter is unable to recognize characters of the alternatives.

The equivalent backward approach actually has a very logical basis: users who start a job using the alternate method are then subject to the keyboard to finish the job. This method can, however, be highly confusing to the novice user. Let's defer examination of the possible confusion until you see what alternatives to the keyboard exist.

When the keyboard furnishes input, most command interpreters use the standard DOS buffered keyboard input function. In DOS, function 0AH (contains raw input) construction handles through several different keys, including a string edit capability, which we examine in more detail later in this chapter. Note that this editing capability, such as it is, is a feature of OS input function, not of COMMAND.COM itself.

Command-line editors such as Edit, editors (ED), and EMU (NT) and also their own Real-String Function, can be configured to edit the current program or just COMMAND.COM that can function 0AH. These programs include DOS/GEM, SIMMIDE, DOSKEY, the command-line editor mentioned with DOS 5.0, and so on. Instead, it provides separate interface INI 205, ANSIPRM, which a program such as COMMAND.COM must explicitly call, is shown in our earlier example.

On the older operating systems, keyboard input to the command interpreter in DOS is not buffered—precise line passes to the interpreter exactly as you type it, except for characters erased with the backspace key.

All sources of interpreters, such with MS-DOS that I have examined, have a size limit of 120 characters for their keyboard input. This limit is imposed by the limit of the I/O port (see Chapter 7), which—however—28 bytes for the command line. Of these 128 bytes, one is taken by the character count and one by the CR character that terminates the input string.

Although it would be possible to extend this limit by a few bytes because the command line is never copied into the I/O, the actual interpreter does so. For simplicity, most provide a maximum 128-byte input buffer, although at least one (4DOS) allows you to configure a larger buffer size.

**Batch Files** In many applications, a relatively complex series of commands must be entered to get the desired action started. Batch files provide the most generally used method for supplying these

knows. As it type the commands into the batch file, and then the entire sequence is pumped into a command interpreter when you invoke the batch file. Command interpreters treat a batch file as a single type of command. Although different interpreters process these files in different ways, they generally do that the interpreter reads the file line at a time, then executes the command, going back to the beginning of the file when it comes back to read the next.

**COMMANDCOM** causes the batch file each time a line is read, and reopens it to read the next line. This makes it possible to use batch files with a single drive system by having a copy of a batch file on diskette so that multiple copies are named the same and have the same content; the command interpreter would never know that the disks were swapped between lines. This action does, unfortunately, also make batch file operation quite slow.

**Batch Enhancers and Compilers** Batch files are so widely used that several firms offer batch language enhancement programs such as BI Batch Enhancer or the Norton Utilities. BI also offers batch language and several PC Magazine utilities such as Wayne McLeod's BATCHMAN. Recently, batch language compilers have become popular as well. In addition to Wenhart Software's BATCHCOMP, I've seen several others: WILDER BY Moore (in August 1990, he published a batch file compiler, *Don't Panic, It's A BATCH*). These compilers turn BAT files into true COM or EXE files.

This is a very interesting process. As you will know, the SET statement in a BAT file alters the master environment supplied in the section "How COMMANDCOM Uses the Environment", but a workstation can't attempt to change the environment by a COM or EXE program results only in a workstation copy of the program's local copy of the environment, which then gets thrown away when the program exits.

This is why the proper semantics of the SET statement is preserved when a BAT file is compiled into a COM file. Simple. The compiled SET statement uses a undocumented DOS *h* after the master environment. This copy of the *h* in a BAT file with a SET statement is compiled with BATCHCOMP into a true COM file and undocumented INTEL Foundation 52b.

We've seen how to force a session to the DOS List Of Lists, which contains at offset 2 the sequence of the first 16 files in a computer. To examine the MFCB, the program can find the master environment. It is a sparse collection of data later on in this section "Other Ways of Locating the Environment". In any case, the population of the file compiler will, for better or worse, produce many programs that rely on undocumented DOS functions and data structures.

## "Losing" Stuffed Commands

Only the command interpreter can obtain input directly from the batch file, it's not possible to provide direct input to your programs by lines typed into any batch file. However, it's possible to use special utilities to stuff input into the keyboard buffer, where your programs can read it, and facilities can be called from the batch file. This is useful with programs that force file command-line arguments.

One such keyboard-stuffer is Charles Petzold's KEYFAKE.COM, available in the book *PC Migration: The DOS Power Tools*. KEYFAKE is used here to illustrate a feature of batch files that, although seemingly obvious, appears because a lot of confusion. Any other keyboard stuffer serves equally well.

The following batch file creates a file called FOO.BAR, containing the single line "hello" (the 13 26 13 emits a newline - 2 newline sequence).

```
echo off
key fake "hello" 13 26 13
copy con foo.bar
```

Note that input is stuffed into the keyboard buffer before the COPY CON command is invoked. If you want to stuff both the input and the command into the keyboard, however, the command must go first:

```
echo off
key-fake "copy con foo.bar" 13 "hello" 13 26 13
```

So far so good. Now let's add a line to the end of the batch file:

```
echo off
key-fake "copy con foo.bar" 13 "hello" 13 26 13
echo Done creating FOO.BAR
```

What happens when you run this? The second command happens before the first command:

```
C:\UNDOC\KYLE>tmp
Done creating FOO.BAR
C:\UNDOC\KYLE>copy con foo.bar
hello
^Z
```

1 File(s) copied

The message that signals that the operation is complete is displayed before the operation begins! Just adding a line to the end of the batch file somehow caused the COPY CON command to be deferred.

It gets worse. Do the keyboard stuffing inside a loop, and the COPY CON command never gets executed, resulting in an infinite loop:

```
echo off
del foo.bar
:loop
key-fake "copy con foo.bar" 13 "hello" 13 26 13
if not exist foo.bar goto loop
```

Simply putting a command inside a loop causes the batch file to stop working! What is going on here?

If, as a final experiment, you return to the original idea of putting the COPY CON command itself in the batch file and stuffing only its input into the keyboard buffer, everything starts working properly again:

```
echo off
del foo.bar
:loop
key-fake "hello" 13 26 13
copy con foo.bar
echo Done creating FOO.BAR
if not exist foo.bar goto loop
```

What you have just seen merely illustrates that the command interpreter exhausts all batch file lines before looking in the buffer for keyboard input. Therefore, unless the keyboard stuffer happens to execute as the last command in the batch file, the stuffed command isn't executed at the correct time.

This detail of command interpreter operation seems rather obvious, yet it spawns at least one question per week on the major network forums that deal with hardware and software problems. The rule to follow in order to avoid the problem is simple: Invoke commands directly from the batch file, not by stuffing the keyboard buffer; provide only program input through the buffer.

## Interpreting Operator Requests

Once the operator's input has been obtained, it must be put into a form acceptable to MS-DOS that will most be parsed for its meaning, correctly interpreted, and acted on. This section first describes what happens inside in parsing the input, then discusses how the input is interpreted, and finally deals with the execution of a terminal command. Any input not recognized as an installed or internal command is passed to the device driver procedure, as a possible external command that it must load from a file.

### Parsing for Inclusion in the PSP

The standardized parse command characters are treated as separators, and only a few of these are used: space, hash (#), dollar (\$), ampersand (&), the switch character (normally a forward slash), and some versions of DOS also can be changed to a hyphen (see below). The comma (the command file character) and the colon (signal) characters have additional syntactic significance, but all are recognized as valid characters in the standard command file name.

In this standardized parse, certain characters are treated as separators, and only a few of these are used: space, hash (#), dollar (\$), ampersand (&), the switch character (normally a forward slash), and some versions of DOS also can be changed to a hyphen (see below). The comma (the command file character) and the colon (signal) characters have additional syntactic significance, but all are recognized as valid characters in the standard command file name.

In the standard parse, a space, or any of the other switch characters at the front of the input line. When a switch character is encountered, the parse routine converts it to uppercase if it is alphabetic and moves it to the beginning of the buffer. From that point on, until the separator characters is encountered, all characters are copied into the parse buffer, and are converted if necessary. When the terminating separator is encountered, the parse pointer is then pointing to:

1. The command in the parse buffer (if it is to be either an installed command or an external program, some version process will occur next). If it is an internal command, the command interpreter will perform the execution. In either case, a PSP is available for the rest of the parsing procedure.

2. A parsing structure from the input buffer are moved to the command tail area of the PSP associated with the command's execution, starting at offset 81h, and the count of those characters (including the terminating CP) is stored at offset 80h. Some of these characters may have been converted during the move.

Next, the first complete word (if any) on the command tail is examined to determine if it could be a filename. It must contain at least one character that would not be valid in a filename, its second character must be a colon, and its second character position (if the word can be a period). This permits the C: drive of DOS 1.x installations (no specification such as FILENAME.TXT) to be accepted.

If the word passes, the test is then applied to the H: drive area, which includes case conversion, corresponding to the H: B area of the PSP at offset 50h (see Chapter 7). The drive code conversion is done by first adding the drive to the hex byte followed by the filename portion padded with spaces to the maximum length. The period (if the colon is omitted) and the extension goes into the extension area, padded with spaces. The process is repeated for the second complete word of a command tail to fill in the H: B2 area of the PSP at offset 64h.

For DOS 1.x, the H: B1 area of the PSP is filled by C: M: drives, and the syntax of many of the other command lines is also done in these positions. In DOS 1.x many programs took advantage of the command line area of the PSP, but DOS 2.x standardized the arguments from the filename fields of H: B1 and H: B2. Some programs were not able to do this, and had to write their own parsing routines, the command file name, and count some of the work for them. However, these routines are not capable of handling switch character references, and for programs (such as DOS 2.x) that wish to begin to take advantage of the command line area of the PSP, unfortunately, some programs still depend on the old way of parsing the command line of the M: entries.

On completion of this standard parse, the command interpreter has in an internal parsing buffer the first word of input converted to uppercase, and it has in the new current PSP the two H: B areas of the command line data. The next step is to determine whether the input was actually a valid command. If it is not, it is known on the subject of the

**SWITCHAR** If you've ever needed to switch between DOS and UNIX machines, you may have been annoyed that whereas UNIX uses the forward slash / for paths and hyphens - for command-line options, MS-DOS uses the backslash \ for paths and the forward slash / for command-line options. Prior to DOS 5.0, the convention applied to COMMAND.COM (more than to DOS itself) since DOS 5.0, though it's built into DOS and cannot be modified.

What a mess. An undocumented DOS function, INT 31h (function 3701h), can help you get up this situation if you use a version prior to DOS 5.0. This function, as described in the Appendix, changes the switch character. This facility was documented for a brief time in the DOS user interface (the SWITCHAR option in the DOS 7.0 CONFIG.SYS) but then it was made undocumented.

This function can be incorporated into a tool utility (see Fig. 10-1) that sets the DOS SWITCHAR. Packages of UNIX utilities for DOS, such as the MKS Toolkit, include a similar utility.

**Figure 10-1. SWITCHAR.C**

```
/*
SWITCHAR.C -- uses undocumented DOS function 3701h
switchar changes DOS switch char to - and path char to /
switchar \ restores DOS switch char to / and path char to \
*/

#include <stdlib.h>
#include <stdio.h>
#include <dos.h>

main(int argc, char *argv[])
{ int c = (argc > 1) ? argv[1][0] : '-';
  if(c != '\\')
    _DL = c;
    _AL = 0x3701;
    geninterrupt(0x21);
    _AH = 0; /* value returned in AX */
  }
  else
    _asm {
      mov dl, c
      mov ax, 3701h
      int 21h
      xor ah, ah /* value returned in AX
    }
}
endif
}
```

Most of COMMAND.COM since DOS 4.0 (including the DIR command) supports the SWITCHAR. Setting SWITCHAR with function 3701h is useful only when other programs in particular COMMAND.COM (which call function 3700h) use SWITCHAR. As of DOS 5.0, the SWITCHAR program won't even work, but I used it on a 386 to set a switch function 3701h (now as expected) to make the program an exercise in portability for DOS 5.0 and above.

**Command Line Redirection and Pipes** One feature that has existed in COMMAND.COM since DOS 2.0 (and which consequently is also in all the popular alternative command interpreters) is the ability to redirect input and output. This means that the command interpreter can "throw a switch" so that input for a process comes from a file rather than from the keyboard and can be piped (i.e., sent) output from a process to a file rather than to the CRT.

Logic or workflow features, which came from UNIX "pipes" can be created to look to a process for input, but the output of the first becomes the input to the second.

The capabilities described here were documented the method by which they are achieved in Fig. 10-2. Let's look first at how COMMAND.COM performs this magic in DOS 5.0. It's to create a simple C program that performs its own redirection, using command-line arguments to specify the input and output names.

Here, `>` shows command line redirection works when COMMAND.COM processes this line:

```
D:\J0052\ISPT> DEV >xyz
```

The redirection sequence specifies that output is to go into the file `XYZ` rather than to the screen. `DEV` is the device driver program from Chapter 7, but any program could have been used. And though this example deals only with redirection of output, output redirection is done in exactly the same way.

COMMAND.COM processes `DEV >XYZ` in almost the same way that it processes a plain `DEV` command. Only after `DEV` has been determined to be an external command, the file located, and the current working directory is name retrieved, does the command line redirection cause a change. First, COMMAND.COM processes `DEV >XYZ`. COMMAND.COM creates a file named `XYZ` in the current working directory.

The System File Table entry for the handle returned by the Create File call (0x121h AX=30 04h) is then opened into `STDOUT` handle 1. All output sent to `STDOUT` from this point on goes into `XYZ` rather than to the CRT. COMMAND.COM then continues just as in the earlier example to load and run the program `DEV`.

After returning control to its process, COMMAND.COM detects that `STDOUT` handle 1 points to a different System File Table entry from `STDERR` handle 2. This indicates that `STDOUT` has been redirected. Consequently, COMMAND.COM closes `STDOUT` 1 and restores it to the same SF Table entry as before, then opens into the loop that closes the 45 interruptible handles. Hence, what most programmers don't know how to treat as `STDERR`.

Because of the device independence achieved by the use of the handles and device drivers (device driver "Comptel" and those two add-ons to the normal inner loop of COMMAND.COM were kept out from the scope of `DEV` into file `XYZ` rather than onto the screen.

The gap between redirection as described here and pipes that connect processes is that the normal inner loop of COMMAND.COM creates a temporary file to serve as a pipe and redirects output into it. The second process then directs the file to the second process as its input. Thus, instead of the normal command interpreter loop that generates the prompt and wait for keyboard input, are stopped for the second process so that the single piped command actually does the same thing as two separate commands set a common redirection file connecting them. Finally, the command interpreter creates a new list over the temporary file when it's no longer needed.

You see, even though COMMAND.COM deals with command line redirection. Now let's see how to do this in a C program (code, Figure 10-7). The key is the `dup2` function.

Figure 10-2. REDIR.C

```
/*
 * REDIR.C - August 1992 - Jim Kyle
 * Shows how to redirect stdin and stdout
 * Tested only with BCC++ in ANSI C mode!
 *      gcc redir.c
 */
#include <stdio.h>
#include <string.h>
#include <io.h>
#include <process.h>

int main( int argc, char **argv )
{
    char cmdbuf[128] = "";
    FILE *NewIn = NULL, *NewOut = NULL,
    int oldIn, oldOut;
    int retVal = 255, i;

    if( argc < 4 )
        puts( " Usage: REDIR In Out Command ... " );
}
```



```

puts( " where in = STDIN or name of input file" ),
puts( "      Out = STDOUT or name of output file" );
puts( " and Command = normal command line to be run" ),
}
else
{
if( strcmp( argv[1], "STDIN" ) != 0 ) /* if file named */
{
NewIn = fopen( argv[1], "r" ); /* open the file to read */
if( !NewIn ) /* quit if any error */
{
perror( argv[1] );
return retval;
}
oldin = fileno( stdin ); /* save for restoring */
dup2( fileno( NewIn ), fileno( stdin ) ); /* force to stdin */
}

if( strcmp( argv[2], "STDOUT" ) != 0 ) /* if file named */
{
NewOut = fopen( argv[2], "w" ) /* open the file, output */
if( !NewOut ) /* quit if any error */
{
perror( argv[2] );
return retval;
}
oldout = fileno( stdout ); /* save for restoring */
dup2( fileno( NewOut ), fileno( stdout ) ); /* force stdout */
}

for ( i=3, iargc, i++ ) /* rebuild the command line */
{
strcat( cmdbfr, argv[i] );
strcat( cmdbfr, " " );
}
retval = system( cmdbfr ); /* then exec, save exit code */

if( NewIn ) /* restore STDIN if needed */
{
fclose( NewIn );
dup2( oldin, fileno( stdin ) ); /* use original value */
}

if( NewOut ) /* and also STDOUT */
{
fclose( NewOut ); /* close to update directory */
dup2( oldout, fileno( stdout ) ); /* use original value */
}
}
return retval; /* return exit code */
}

```

Rather than cause a syntax error by attempting to parse its input command line, to detect a redirection request REDDIR just requires that you specify both an input and output path and the correct two command-line arguments. If the first argument is "STDIN" no input redirection is performed. If it is anything else, an attempt is passed as a file or device name to the `fopen()` function to be opened for reading.

Similarly, the second argument is compared to "STDOUT" to detect a request for output redirection.

After these two arguments are processed, the rest of the input command line is assembled into a new command line and passed on to the normal command interpreter through the `system()` function. Any value returned by `system()` is saved in the local variable `retval` to be passed back as REDDIR's own exit code. We look at this exit code business a little later.

Before returning to DDBS REDDIR closes any redirection files or devices that have been opened and attempts to restore the corresponding stream to its original device. With these tasks taken care of, REDDIR then returns the saved exit code as its own and goes back to DDBS.

The `dup2()` function is a standard C library function whose MS-DOS implementation is documented in 1.21b: AH 46h: Force Duplicate File Handle.

**Distinguishing Internal and External Commands** Every command interpreter implemented for a microcomputer has included at least some internal commands; most have also provided a means of executing external commands.

What distinguishes an internal from an external command is the location of the code that executes it. All commands, whether internal or external, reside in the command interpreter itself; external commands reside elsewhere. In some systems, including Tandem's TRSDOS, some microprocessors, and Honeywells' G/OS, the external commands, external commands have been stored in special library files. In MS-DOS, however, they are stored as individual program files.

CP/M 2.2, when the MS-DOS architecture was derived, contained only five internal commands. All other commands were external, stored as COM or COMMANd files in a memory image format. The original commands were DIR, REN, COPY, ERASE, and SAVE. The first four of these were functionally identical to those in MS-DOS (see below). The other one, provide, I like, means for creating the necessary COM files. To avoid SAVE, the specified number of 256-byte "pages" to a named file. To find a really useful program, even RAM, so that SAVE could do its thing, you had to use DD.COM (available in the Best of DOS debugger).

The earliest version of MS-DOS had only a few more in the way of internal commands: COPY, CHECK, and C.M. had been functions of the external utility called PIP (Peripheral Interchange Program) in the early retail DOS systems, moved into the command interpreter as an internal, and DATE, TIME, DIR, and CHDIR added. With each new version, as features were added to the system, additional external commands came with them. The use of batch files gave birth to several internal commands to help solve some of system trouble.

By the time MS-DOS made it to version 3.3, the list of internal commands had grown to 36. One undocumented command was added with version 4.0, two documented commands at version 5.0, and none in version 6.0.

The undocumented TRUENAME command displays the full path name of a file, gives the file's name as its argument. An MSDOS implementation is translated from logical back to physical form, and an explicit path is not explicit. For instance, on one system the command SUBST C:\ZAP\ZIP and C:\ZAP\ZIP\ZSSC. This command corresponds exactly to INE2H Function 60h, which is explained in greater detail in Chapter 8 on the DOS file system. Operation of the new COMMAND and CHDIR commands is covered in "Comparing" in the discussion of the HMA and UMBs.

Let's get back to the way that a command interpreter determines whether it's dealing with an internal command or an external one. Don't make a single search of its internal command list. If the input command is found, it has its internal status list, its internal and the corresponding external command names, and the file is not found, the interpreter treats this as a possible external command.

DOS 3.3 introduced a process of searching the internal command list by way of INRs that contained the internal COMMAND.COM names, a set of undocumented tools. However, because this feature was not published, external commands are produced with the exception of APPEND if DOS's standard setup program, which was created by the first place, and DOSSHELL, which appeared first in DOS 4.0. Some of the new programs exhibited serious problems, and in DOS 5.0 the original DOSSHELL, supporting two parallel program packages under the same name, created by Central Point Software, replaced the previous one. We provide source code for an installable command list in this chapter, but you can get the books that MS-DOS provides for command interpreters.

Note that the file name command that is passed to the interpreter is parsed before any attempt is made to locate the command. The percent character and the interpreter treats the percent character as having special meaning because of its use in references to command-line arguments when found in a batch file; it also identifies environment variables. If followed by a character that is neither an argument identifier nor an environment variable, the percent character normally is thrown away

rather than pass `FILE` to the command as part of the command line. This is true only when input is taken from a batch file.

Sometimes the character needs to be kept, though. The classic case is the `FOR` command with its inherent variable reference, as in `FOR %I IN * DO type %I`. This one works perfectly from key word input, but if it is included in a batch file, `%` characters will be dropped, and the command line will generate a syntax error. To solve this problem, whenever `COMMAND.COM` or its function equivalent creates a new period character, it replaces them with `%%` signs, as with `COMMAND.COM` versus `FILE` in a program. Thus, a batch file requires `"FOR %I IN * DO type %I"`. This notion is analogous to the C++ language conventions regarding `\\`, which requires that you use `\\` in any string where you want a single `\` to appear.

**Finding and Executing Internal Commands** Because the internal command list is using an installed command's search path, it's a tricky matter to copy `COMMAND.COM` to run a program that uses the same name as one of the internal commands. That is, you name a program file `TYPE.COM`, you find that the internal command `TYPE` takes its place when you try to execute the program, as if the command interpreter, through it, runs your own program perfectly when invoked through the `HOSTENT` function `TYPE`.

When `COMMAND.COM` searches its command list for a possible internal command, it breaks off the command word to search its installed period after the command. Thus, to continue the example of your named `TYPE.COM` that you copy, the `COM` would be ignored, and the internal search would find and match the `TYPE`, which it would then execute.

With more recent versions of DOS that permit you to specify a full path name for a command, you can solve this problem by specifying the full path to `TYPE.COM`, such as `TYPE`, if the file is in the current directory.

A vast amount of the alternative command replacements attack this problem in two ways. First, they let you selectively disable a particular command having a configuration option second lines vary the way these search paths perform. The second way is simply that the input is not truncated at a period. If you name `TYPE.COM` that is similar to a system file, and of course it will run, find `TYPE.COM` because none of the internal commands have a period in their names.

Under an alternative replacement, once an internal command is detected, it is normally executed inside that file. The user is not allowed to see the command list or the appropriate internal command followed by a loop back to the prompt code, because none of this code is constructed within the interpreter, so the program swapping occurs and the child process is spawned for the execution of internal commands.

Before we alter a command `COM`, as introduced, the external command `COMMAND` with the `COPY` or switch is used to remove batch file from inside another, then return to the original file. This technique still works, but requires about 4k additional RAM for the child copy of `COMMAND.COM` that is loaded, as well as, may be significantly slower due to the additional disk accesses required.

### Dispatching Appropriate Processes

Any name that is not an internal command or matched in the installed command list is presumed to be an external command. The command interpreter searches for a file of that name, using the `PATH` to determine where to search. If such a file is located or is loaded and run from the error message, "Bad or missing file name," tells the operator that the input was faulty, the command interpreter then returns to its top-level "get input" procedure for a fresh command to interpret.

**Locating and Loading External Commands** To locate and execute an external command, the command interpreter first searches the current working directory for a file having a name identical to the command word and the extension `.COM`. If this fails, the search is repeated in the

extension .EXE, if this also fails, a third search is made with the extension .BAT. Thus, if three files named `FOO.ME`, `COM.FOO.ME`, `EXE`, and `FOO.ME.BAT` all exist in the current directory, only the `COM` file will be executed as an external command.

The `COM`, `EXE`, `BAT` sequence is established by code in the command interpreter, not by MS-DOS itself. It can be explored in several ways that we examine a little later, but if you want to change, or at any rate, control, your behavior by locating in your command interpreter the nine bytes that contain these characters `COMEXE.BAT`, it is changing them as you desire. This can be done on the disk copy of the interpreter or in memory using `DEBUG`.

If all three searches in the current working directory, the command interpreter looks for an environmental variable named `PATH`, and if one exists, it takes the first path listed in that string (that is, all characters up to, but not including, the leftmost semicolon) and repeats the triple search in the directory specified by that path. If no path exists, the interpreter moves on to the next pathspec in the `PATH` variable and repeats the search. This continues until either it finds a file that satisfies the search, or the `PATH` variable is exhausted without finding such a file.

If a file is found, the command interpreter issues a "Bad filename or command" error message in Function key prompt. Note that this happens only after the three searches `COM`, `EXE`, `BAT` have been made in each directory specified by `PATH`. If you have a large number of directories in your `PATH` and if each has a large number of files, the search may take a significant amount of time.

**Dealing with BAT Files** When a file is found that satisfies the search criteria, the command interpreter must determine upon whether or not the file found was a batch file (BAT extension). If so, the command interpreter sets appropriate flags to indicate to itself that it is processing a batch file rather than a regular program; the interpreter also stores enough data about the file to be able to find it again. The flag locations and the amount of data saved vary significantly from one version of DOS to the next.

The interpreter opens the batch file, reads its first line into the command input buffer, and replaces any indicated arguments (e.g., `%1`, `%2`, and so on) with corresponding words from the original input line, since a normal program expects a buffer. If any `%` argument is found, it is replaced with the name of the original file, suitably extended. The interpreter then closes the batch file and interprets and executes the modified batch line, just as if that line was typed from the keyboard.

When execution is completed, control returns to the command interpreter. The flags tell the command interpreter to interpret the next line, read the next line (subject to a minimum of 32 bytes) into the buffer. The process repeats until all lines of the batch file have been executed.

The command interpreter can be used to execute either internal or external, and since DOS 3.3, external and internal, as well as additional, so-called, batch files in a subprogram fashion using the `CALL` command. For any reason, it is entirely possible to invoke a batch file that never finishes executing, or, at least, a batch file never gets back to the original command interpreter's keyboard input level.

**Dealing with COM and EXE Files** If the file found was not a batch file, the command interpreter uses the `DOS.EXE` function (INT 21H, function 4B00H) to spawn the file's execution as a child process of the command interpreter. In DOS 3.0, this happens as the interpreter until the child process terminates.

Notice that with `COMMAND.COM` it makes no difference whether a `COM` or `EXE` file was found. The distinction between the two opposite executable files is made by the `EXEC` function based on one or the other two bytes of the file's file extension. `EXE` files have the MZ signature, and the actual file extension is ignored, except to make the two files mean that you could force a file that is really an `EXE` type to be found during the first search by changing its extension to `COM`. It also means that you could take a text file (word processor document, spreadsheet, etc.) and give it with an extension of `COM` or `EXE`, and `COMMAND.COM` would try to execute it as a `COM` file.

An alternate to changing the file's extension is to create a stub loader, which is a small file of the same name, except for the extension COM. The stub loader then uses the EXEC function to spawn the original EXE file as a child of its own. This approach makes it possible to set up all sorts of special conditions before the program file is executed, then subsequently undo them with minimal overhead.

One widespread use of this technique is in support of the third-party replacement video BIOS package UltraVision, from Permonics. UltraVision permits its users to set up a wide variety of screen options, ranging up to 132x60. These formats are compatible with any other program that is well behaved (in the sense that it looks in the BIOS work area to determine what number of columns and rows are currently set). Unfortunately, many popular programs don't bother to do so because the ability to set non-standard formats is relatively recent.

For any programs that do not check the current number of columns and rows, it is necessary to set the video format to the standard 80x25 dimensions. Permonics UltraVision includes a utility which generates a stub loader for any desired EXE-format program file. The loader first notes all pertinent information from the BIOS area and then resets to the 80x25 format. Next it EXECs the specified program file, using its full name (including the EXE extension) and passes to it all the command-line arguments that the loader itself received. Upon return, the loader temporarily saves the exit code from the real program, stores the video set up to the conditions the loader found at entry, and returns its child's exit code to DOS as though the code were its own. The command interpreter always hands the loader with the COM extension first, so as to ease conversion to a .com file.

As a result, by using the right-supplied utility UltraVision to create a loader, you can make any program well-behaved in the video area at the cost of a few thousand bytes of overhead code.

Another use of the stub loader idea is built into Microsoft's new segmented executable files which are used in Windows OS/2 and are now obsolete. European OEMs manufacturing IBM PC-like PCs use EXE files with a normal 44-byte EXE file at their head, followed by a 2-MB header with the NT operating system. The old-style MS-DOS loader is programmed to print a message, such as "This program requires Microsoft Windows", if it can be used as a loader that, for example, runs Windows.

**The Exitcode Idea** When I said that the UltraVision loader saved the real program's exit code, you may have wondered why a loader should preserve the exit code of a spawned process, or possibly even what an exit code is all about. Although it is documented at least with regard to the method of saving a program's exit code to its parent, the details are so spread out that some explanation is in order.

To recall that a process should retrieve a result code to its parent, it is generally followed directly from the idea that every process in a system is, in a sense, a child of its parent, that is, each has some higher-level process of which it is a child. In our case, the parent is the loader, thus the apparent complication of about the same time as the idea of the operating system itself, only before the operating system came upon the scene.

A DOS process can retrieve exit code to its parent. The exit code is returned to the program if the process terminates using INT 2xh functions 4Ch (Terminate with Exit Code) or 41h (Terminate But Not Respond). If any other method is used to return to DOS, such as INT 20h or INT 21h (Exit on Error), the exit code is always set to zero as generated by DOS itself.

The code is not retrieved once and only once after the spawned process returns to the parent and before any subsequent process is spawned; the reason for this restriction is that DOS itself provides only one location to hold the exit code, so that every process overwrites the code left there by its predecessor (if any). The BIOS function that retrieves the code (INT 21h (Exit on 4Dh) Get Exit Code) is a programmatic error in that storage location in the process of retrieving the code.

COMMAND.COM provides the internal command ERRORLEVEL, which is actually a function that can be accessed by the internal IF command. ERRORLEVEL retrieves the exit code of the most recent command and compares it to a specified value. All command parameters and switches provide this command with its implementation in a functionally identical manner.

It is a little tricky to see how the `COMMAND.COM`'s processing loop that began this chapter, you saw that the first thing `COMMAND.COM` does upon return from a child process is to retrieve the exit code and save it in `COMMAND.COM`'s own data area for possible use by `ERRORLEVEL` later. When `ERRORLEVEL` is asked, it uses that saved value rather than attempt to retrieve the exit code from DOS again.

The idea of exit code and `ERRORLEVEL` applies only to external commands; most internal commands have one (located) on the `ERRORLEVEL` value. Because these commands do not spawn child processes, these commands do not affect the DOS exit code value. Some third-party command interpreters do enhance the `ERRORLEVEL` idea to at least some internal commands, extending it to all commands, including those that test it would negate the whole idea because it would make multiple way decisions impossible.

### Installable Commands

In action `ME` is installable. Command `0x1212h` available since DOS 3.3 connects what would ordinarily be an external command into the internal command list of `COMMAND.COM`. Both the command word and the handler for this instruction must be installed as a single, non-pop-up ISR. The original instruction was expected to be an `APPEND` in DOS 3.3 and higher and in the DOS 4.0 `DOSSETUP` facility, which is now of only historical interest since the DOS 5.0 version of `DOSSETUP` is totally different.

You can use the undocumented interrupt to install your own commands with DOS 3.3 and later. A new task pointer, "Does anyone see I call you?" `COMMAND.COM` issues calls to this interrupt if you place a file pointing to a buffer. To use the functionality you provide your own interrupt handler, hooking `INT 12h` `AH=Afh`.

Each time `COMMAND.COM` saves its Subfunction `AH00h` are to determine if the command is not installed as external command. The other two are called only if the first call returns `FFh` in the `AL` register, indicating that the command is indeed valid. The second call is both `CS` and `SI` in `AH01h`, which specifies to execute the command.

When Subfunction `AH00` is called the first time, both the `DS` and `ES` registers point to `COMMAND.COM`'s memory space. The `DX` register is set to `0011h` for unknown reasons, the `BX` register points to a buffer that contains two copies of the command followed by an exact copy of the input line. `CH` contains `1Fh` and `AL` register points to a buffer of buffer that contains only the command word from the input line. It is supposed to be used and processed by its character count. Your code must verify that the input of this command buffer is a valid handle for the name of your installed command. If it is not a valid one, you should leave the interrupt group the one in case some other installed command is going to be called. If it is a valid one, the address must come installed routine, your routine should change the `AL` register to `Fh` and return to the user.

The second call to Subfunction `AH00h` is identical to the first except that `CH` contains `00h` and `SI` points to `COMMAND.COM`'s code under determine why `COMMAND.COM` makes both calls. The alternate call is used only if the `COMMAND.COM` does it a satisfactory. It calls this subfunction only once.

The Subfunction `AH00` returns `00h` in `AL` indicating that the command is not an installed command and `FFh` in `AL` if it is a valid. Instead, the command interpreter is expected to deal with the command itself.

If Subfunction `AH01` is called to execute the command, the `ES`, `DS`, `BX`, `DX`, and `SI` registers contain exactly the same data as in the previous subfunction, although only `DX` and `SI` are again explicitly loaded with the values. Due to the call, the `BX` register has been preserved through a couple of subroutines calls, but in future DOS versions may not always retain the pointer to the full input line buffer.

Therefore, if your command accepts arguments on the input line, we recommend that those arguments be copied into a local buffer before the routine returns from the AEO0h call. After executing the actual command within the AEO0h call, your code should zero out the contents of the command buffer in DOS. The zero-out signals COMMAND.COM to ignore, not to execute, the command itself.

These calls to Function AEOh are made before the original command is executed, which means that you can install a command that has the same name as one of the original commands and your command will replace the original one. Because your own commands do nothing, this offers an elegant way to disable the DEF and ERASE commands, or a system that can be accessible to the general public.

This also permits you to add a modified function to any original command. You can provide an installed command with an identical name that does your local command processing (such as call to Subroutine AEO0h) and sets non-zero on the contents of the command buffer, and returns zero from the AEO0h so that COMMAND.COM will process the original command.

The C program shown in Figure 10-3, INSTCMD.C, illustrates all the points of dealing with the installable command intercept service. Note that its program does not install as a TSR but instead operates as a wrapper, see "Application Wrappers for DOS" *Technique* in *July 1992* to avoid problems associated with the added command. Note that, despite the reference to COMMAND.COM in the comments, this program actually works with any command interpreter that supports 21. At some, the program uses the COMSPEC environment variable to locate the command interpreter itself.

**Figure 10-3. INSTCMD.C**

```
/*
INSTCMD.C
```

The installable Command function is not called by a program that wants to extend COMMAND.COM's repertoire. Instead, you hook the function and wait for COMMAND.COM to call you (don't call us, we'll call you). Function AEO0h lets you tell COMMAND.COM whether you want to handle the command, and function AEO1h is where you actually handle it (similar to device driver division of labor between Strategy and Interrupt).

Note that AEO1h is called with only the name of the command; not with any arguments. Therefore, arguments must be saved away in AEO0h. You'

Furthermore, while redirection is handled in the normal way in AEO1, in AEO0 we get the entire command string, including any redirection or piping. Therefore, these must be stripped off before saving away the args during AEO0 processing.

Problem with the following AEO0 and AEO1 handlers: they should chain to previous handler. For example, (INTRSPY program won't see AEO0 and AEO1 once INSTCMD is installed).

The sample COMMAND.COM extension used here is FULLNAME, based on the undocumented FRJENAME command in DOS 4+. We simply run undocumented Function 60h in order to provide FULLNAME. Actually, not quite so simple, since function 60h doesn't like leading or trailing spaces. These are handled inside function fullname().

The following INTRSPY script was helpful in debugging ZFAE.

```
; INSTCMD SCR
structure cmdLine fields
max (byte)
```

```

text (byte,string,64)

intercept 2fh
function Dash
  subfunction 00h
    on_entry
      If (dx == 0FFFFh)
        output "AE00"
        output (DS BX->newline)
        output CH ; CH=FF (first), 0 (second)
        output " "
      subfunction 01h
        on_entry
          If (dx == 0FFFFh)
            output "AE01"
            output (DS-SI->byte,string,64)

tested only for Microsoft C 6.0+ or above
cl -qs instcmd.c
*/

#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <ctype.h>
#include <conio.h>
#include <dos.h>

#pragma pack(1)

typedef struct {
  unsigned es,ds,d1,d2,bp,sp,bx,dx,cx,ax;
  unsigned ip,cs,flags,
} REG_PARAMS,

typedef unsigned char BYTE;

typedef struct {
  BYTE len;
  BYTE txt[1];
} STRING,

typedef struct {
  BYTE max;
  STRING s;
} CMDLINE,

void interrupt far handler_2f(REG_PARAMS r),
void (interrupt far *old)();
void fail(char *s) { puts(s), exit(1); }

main(void)
{ /* hook INT 2F */
  old = _dos_getvect(0x2f),
  _dos_setvect(0x2f, handler_2f),
  puts("This demo of installable commands isn't a TSR.");
  puts("Instead, it just creates a subshell, from which you can EXIT.");
  puts("When done, in the subshell, one new command has been added:");
  puts("FULLNAME [filename].");

  system (getenv( "CONSPEC" ));

```



```

/* unhook INT 2f */
_dos_setvect(0x2f, old),
)

char far *fullname(char far *s, char far *d)
{
    char far *s2;

    /* INT 2fh AH=60h doesn't like leading or trailing blanks */
    while (!isspace(*s))
        s++;
    s2 = s;
    while (*s2) s2++;
    s2--;
    while (isspace(*s2))
        *s2-- = 0;

    __asm {
        push di
        push si
        lea di, d
        lds si, s
        mov ah, 60h
        int 2fh
        pop si
        pop di
    }

    return d;
error:
    return (char far *) 0;
}

void fputs(char far *s)
{
    /* can't use stdio (e.g., putchar) inside 2FAE01 handler */
    while (*s)
        putchar(*s++),
        putchar(0x0d);
}

char buf[128], /* not reentrant */
char args[128],

#define CMD_LEN 8

void interrupt far handler_2f(REG_PARAMS r)
{
    if ((r.es == 0xA00) && (r.di == 0xFFFF))
    {
        CMDLINE far *cmdline,
        int len;
        FP_SEG(cmdline) = r.ds,
        FP_OFF(cmdline) = r.bx,
        len = min(CMD_LEN, cmdline->s.len);
        if (!memcmp(cmdline->s.txt, "fullname", len) == 0) ||
            (!memcmp(cmdline->s.txt, "FULLNAME", len) == 0))
        {
            char far *redir,
            int argslen = cmdline->s.len - CMD_LEN;
            memcpy(args, cmdline->s.txt + CMD_LEN, argslen),
            args[argslen] = 0,
            /* yuk! we have to get rid of redirection ourselves! */
            /* it will still take effect in AEO1 */
            /* the following is not really correct, but okay for now */
            if (redir = _fstrchr(args, '>'))
                *redir = 0;
        }
    }
}

```

```

    if (redir == _fstrchr(args, '<'))
        *redir = 0;
    if (redir == _fstrchr(args, '|'))
        *redir = 0;
    r_ax = 0xAEEF; /* we will handle this one! */
}
else if ((r_ax == 0xAED1) && (r_dx == 0xFFFF))
{
    STRING far *s;
    int len;
    FP_SEG(s) = r_ds;
    FP_OFF(s) = r_si;
    len = min(CMD_LEN, s->Len);
    if ((_fmemcmp(s->txt, "fullame", len) == 0) ||
        (_fmemcmp(s->txt, "FULLNAME", len) == 0))
    {
        char far *d;
        if (!*args)
            d = "syntax: fullame [filename]",
        else if (id = fullame(args, buff) == 0)
            d = "invalid filename",
        fputs(d),
        s->er = 0; /* we handled it, COMMAND.COM shouldn't */
    }
}
else
    _chain_intr(nlg);
}

```

If you're an INSDIAD user, first see the four-line explanation of what the program does, followed by the four-line assembly. In a copy of your command interpreter, you'll find the added command FULLNAME. This command is available only if you're using the undocumented TRUENAME command introduced with DOS 4.0 and described elsewhere in this chapter.

Next, in the INSDIAD, we copy the arguments supplied on the input line into a buffer when working with the MSDOS then we use that when executing the command subsequently. Both arguments are in the function handler. 2b

For another view of internal DOS commands, see John Prowse's article "Replacing Internal DOS Commands" (*PC World*, July 1991) and several other articles on 21 AF cited in the bibliography.

### ***TSHLL, a Simple Command Interpreter***

The same is expected of any command interpreter: it's extremely simple. What makes them seem complicated is that they're made as flexible as possible, each instance with a minimum disruption of system operation. To do that, we use a special command interpreter can be a standard function, Figure 10-4 shows TSHLL.C and its execution on a standard as the primary shell of your system.

**Figure 10-4. TSHLL.C**

```

/*****
 * TSHLL.C - Demonstration tiny command interpreter
 * Jim Kyle, July 10, 1990
 *
 * Intended only to show basic principles, not for use
 * with DOS versions prior to 3.1 (EXEC function of such
 * versions does not preserve stack registers).
 *
 * For Borland C compilers only due to pseudovariables
 *****/

```

```

*      gcc -mt -c tshell                                *
*      link /t /c c0t\tshell,tshell,,cs.lib            *
*
*****;
#include <stdio.h>
#include <string.h>
#include <dos.h>
#include <dir.h>

char cmdbuf[128];
char *cmdlst[] = {"DIR", "RUN"},
int i;

void do_dir( void )                                /* reports files in cur dir */
{ struct fblk wkarea,
  int endir;

  if (strlen(cmdbuf) < 5)                          /* default to all files */
    strcpy( cmdbuf+4, "- *.*" );
  puts( "\n Files and sizes\n" );
  endir = findfirst( cmdbuf+4, &wkarea, 0 );
  while ( !endir )
    ( printf( "%-15s %8ld\n",
      wkarea.ff_name, wkarea.ff_size );
    endir = findnext( &wkarea );
  )
  putchar( '\n' );
}

/* CHEAT! Implemented RUN by EXECing COMMAND.COM. */
void do_run( void )                                /* caution, safe only for DOS! */
{ struct {
  unsigned eseg, clo, cls;
  long fcb1, fcb2;
} parms;

  cmdbuf[0] = strlen( cmdbuf+1), /* was wrong in first edition */
  cmdbuf[1] = '?',
  cmdbuf[2] = 'C',
  cmdbuf[3] = ' ';
  parms.eseg = 0,
  parms.clo = (unsigned) cmdbuf, /* was wrong in first edition */
  parms.cls = _DS,
  parms.fcb1 = parms.fcb2 = 0L,
  _ES = _SS;
  _BX = (unsigned) &parms;
  _DX = (unsigned) C:\COMMAND.COM; /* may need to change this */
  _AX = 0x4B0D;
  geninterrupt( 0x21 );
}

void main( void )
{ puts( " TINY SHELL DEMONSTRATOR\n" );
  puts( "Copyright 1990 by Jim Kyle\n" );
  puts( " Commands: DIR, RUN only\n" );
  for( ; ; )
    ( printf( "tinyshell> ",
      gets(cmdbuf);
      for(i=0; i<2; i++)
        if(strlen(cmp(cmdlst[i],cmdbuf,strlen(cmdlst[i]))==0)
          break;
      switch(i)
      { case 0: do_dir();
        break;
        case 1: do_run();

```

```

break,
default: puts("Unknown command!\n"),
}
}

```

Of course, this shows just the bare rudiments of a DOS command interpreter. A real one would have to handle INI 23b Ctrl-C and INI 24b Critical Error, at the very minimum. It also would be important to take over enforcement of INI 25b discussed later in this chapter, even if only to point it at an INT 1 instruction. While it might be possible to stumble through these things at a bare bones level, as do the kernels, even if one barely knows what's totally obscure, the point being made, COMMANDCOM for DOS 8.0, which disassembled, generates 38K pages of listing, at least half of that hardly errors!

Unlike the other programs in this chapter, ISHLL.C is written to be compiled with only Borland C++ or Turbo C. The interface to the DOS EXEC function in this program cannot rely on an environment being established, remember, COMMAND.COM isn't running, and the binary functions of both the Microsoft and Borland products do not do this.

The communication between the command interpreter and the environment variables can be misleadingly discussed as a matter of convenience. The normal command interpreters allow you to modify the prompt and to specify the path through which any program can spawn a copy of the interpreter. Both functions can be accomplished by setting ASCII text strings, environment variables named PROMPT and COMSPEC. Similarly, you can specify the paths to be searched for external commands, and they are stored in the PATH environment variable. COMMAND.COM and other command interpreters therefore have information from these variables in the environment. But as shown by ISHLL, there's no real, or any necessary, connection between the command interpreter and the environment, it's all a matter of convenience and convention.

Note that ISHLL is not intended to be used. Its sole purpose is to present the skeleton of a command interpreter. It recognizes only two commands, DIR and RUN. The RUN command actually launches COMMAND.COM as a child process in order to make the conventional DOS command set available. Without this capability, it's much more difficult to get ISHLL out of your system once you've tried it out!

Once you have ISHLL.COM added to, such as the following to your CONFIG.SYS file (temporarily REM out any existing SHELL = statement):

```
SHELL=C:\TSHLL.COM
```

Note that you must use the SHELL = command in CONFIG.SYS to install a different command interpreter. Many people are under the impression that a new shell can be so easily installed by changing the COMSPEC = variable in the environment. As the environment variable is used only after a command or to process a sub-command, the child command interpreter is loaded and initialized, neither the COMSPEC variable nor the master environment block itself exists!

Now, to see the system. You'll see that the prompt is now C:\Tshell> rather than anything you may be used to seeing. You'll see also that AUTOEXEC.VBI did not run and none of your TSR programs have been installed.

If you type DIR, you get a list of files in the current directory, but no subdirectories will be shown, not with the amount of available space. Any other conventional command produces only an error message because ISHLL does not recognize it.

When you type RUN, you have all normal commands available to you, which lets you run an editor to change CONFIG.SYS back to get rid of the SHELL = ISHLL line.

ISHLL does not create a master environment. You can verify the absence of a master environment by executing the EPSTEAVE program presented here in this chapter. From the child COMMAND.COM spawned by ISHLL's RUN, it shows a master environment, and SET shows that

PATHE is back to the current copy. The lack of an environment's simulation in order to demonstrate the independence of the interpreter from the environment is principle.

The key point about SHELL, though, is the skeletal structure of a command interpreter provided in name. This structure includes an endless loop—direct, minus a prompt, a more complete implementation that uses the environment code interpret and print gets "PROMPT"—gets commands, interprets them, and executes them. That's all a command interpreter does.

## How COMMAND.COM Works

This section is based on the disassembly of several versions of COMMAND.COM; the major emphasis is on those portions already adequately covered in official system documentation. A big section on shell terms is master commitment and primary shell.

Although COMMAND.COM is not the only possible command interpreter, later in this chapter we look briefly at some alternatives. It is the one most used with MS-DOS because, by virtue of part of the system package, it is the operating system because the user system files are loaded into view. Actual to IBM's version of DOS 2.x, COMMAND.COM was an essential part of the operating system; the DOS 3.0 Function-High was contained in COMMAND.COM rather than in the DOS kernel. By the time the Microsoft-IBM versions were released, however, EXEC had been moved to its proper location and the operating system itself became independent of COMMAND.COM.

The process of creating the primary shell actually begins when loading the DOS SYSTEM or HIMEM.COM in PC-DOS loads and its initialization code takes over. After a brief preliminary check on the DOS SYSTEM, we move the program itself that has not yet been executed up to the top-most memory area physically present in the 640K DOS memory. The process is something like the DIBLOAD program in Graphical User Interface during execution.

From this vantage point, most of 386x, the DOS SYSTEM code takes the rest of the DOS kernel, a way that possible for the primary shell to use the DOS services. It is loaded for the first time. In DOS 3.0, DOS SYSTEM also starts to take PSP for its upper memory copy, making possible the use of the 640K-based DOS functions within the initial run code of device drivers, despite warnings of the documentation that it is not the only number of DOS functions and a few others are not.

It is thanks to Hans Sauerberg, creator of BOOBY SYSTEM for venturing ideas. It is strange name though it is appropriate to memory in this case because the Microsoft Chapter 7 is not yet been built.

The CONFIG.SYS file defines the primary shell details, and COMMAND.COM in SHELL. It is necessary in CONFIG.SYS to define the primary shell is essentially the initialization code by DOS SYSTEM after the DOS initialization routine ends the MSDOS.SYS code. The only its initialization routine use INTC2.D function-4000 to load and execute the primary shell program. It is possible to load more than one copies of the command interpreter, some don't follow the COMMAND.COM standard, pointing to itself is its own parent using the PSP. Very soon now we'll take you to the shell's kernel, however who this is done.

One of the parameters passed to DOS with this support is the address of the associated memory block. It is the primary shell's master commitment block just above the DOS kernel area's 256K bytes. Access to the COMMAND.COM initialization routines. The size of this block details to 400 bytes. It is also expanded by adding the /F option variable to SHELL.COMMAND.COM/P line in CONFIG.SYS.

It is particularly interesting that COMMAND.COM, the exact method can differ, but it provide methods for loading files in the master environment to whatever you need.

at version 6.0, DOS prior to 6.0, if the primary shell exercised control to the DOS SYSINITIALIZATION routine, such that a fatal error, such as the stack becoming corrupted or critical system files being overwritten, occurred and that continued operation would be impossible. In such a case, the FDISK utility would be followed by a dynamic boot (IMP 5) which would lock the system down. One possibility, if present, is powering down, and then back up again, so you could return the system to operation.

The primary DOS EXEC file code could gain control was when the primary shell was loading. Next, if the program loaded, could cause CONSOLE SYS INIT to move the program to some directory (probably the root) so that CONSOLE SYS could cause this failure. If this happened with DOS 6.0 or before, you needed to have a bootable floppy handy to bring the system back up so that you could correct the errors in CONSOLE SYS.

With version 6.0, however, DOS INIT is replaced by a prompt that asks for the name of the command interpreter. The program then tries again to load using the name you supply. If that attempt fails, the prompt appears again. This process continues until you give up and use a bootable floppy or until you succeed in establishing a primary shell.

A second possibility is that, during execution, the area at the top of DOS RAM from which BIOS calls the real mode becomes available for reuse. Subsequent operations while initializing the process of booting overwrites the EXEC call and the IMP 5 which follows it.

### Why Shells Are Their Own Parents

It is a common knowledge fact that COMMAND.COM is always its own parent. True, but I wonder how a shell COMMAND.COM is always its own parent, for the very good reason that I can't find any Norton or other 3<sup>rd</sup> edition updated DRIFKIND posted an excellent explanation on the FDISK disk. What is the reason for this message:

Command.com loaded successfully. COMMAND.COM and all other command interpreter shells are designed to run on computers with accepted standards. Since the program contains the "CHECK FOR SHELTER" Abort key priority, Abort, Retry, Ignore, and so on. One option this has is "Help" - a program that has been which means that you are executing an external command that process is immediately terminated.

True, I can't find the code, but I see why executing an external DOS command. The program always reports COMMAND.COM, whatever shell is running, because the shell must contain the "CHECK FOR SHELTER" IN 24 mode, which knows that COMMAND.COM is running and does something other than abort if you press A, right?

When you press the abort key priority, does nothing different, regardless of the process that is running, you press A, the number key priority code of 2, and DOS terminates the current process, and the Shell for COMMAND.COM goes away, changing the system.

The parent process, DOS, always reports the parent PNP field in that process's PNP field. When you press the abort key priority, the PNP field determines which process is going to get control. The current PNP field is the current PNP, however, DOS does not deallocate the program, because it is still active. COMMAND.COM and all compatible shell programs sets the parent process to itself, so that the termination address back into itself. The result is that the DOS parent of the process is the current program, active and retains control.

And, at every COMMAND.COM, the PNP field points to itself. My grateful thanks to D. Riffkind for clearing this up!

### How and Why COMMAND.COM Reloads Itself

One of the least understood parts of COMMAND.COM's internal operations is the reloading of the program portion of COMMAND.COM, upon execution of an external command. If something goes wrong in this process, the error messages can be as confusing as "Bad or missing command interpreter" when it was working fine just a second or two before, to downright misleading ("Unable to

load COMMAND.COM when the real problem was that the COMSPEC variable pointed to a different version of the program. And a point that's usually unclear is why the reloading happens sometimes, but not always.

Whenever an external command is loaded, the RAM occupied by a transient portion of COMMAND.COM is made available to that external command by its use. I decided upon return from an external command, the resident portion of COMMAND.COM does a checkup of the transient area to detect any changes in it.

If any change has occurred, the resident routine records the entire transient area from disk, using the COMSPEC environment variable to locate COMMAND.COM. Only the upper part of the file is loaded. The exact offset into the file at which the record begins is hard-coded into the code that performs the reloading; this offset varies from one DOS version to the next.

This checkup of the transient portion is just calculated immediately after the use of high RAM. That result is stored in the resident portion. From time to time, the checksum is calculated and the result is compared to the original value. A mismatch causes the transient area to be reloaded. Thus, it's essential that the copy of COMMAND.COM pointed to with the COMSPEC variable be identical byte for byte to the copy used at boot time.

After the transient area is loaded from disk, the checksum is recalculated to verify that the reload was in fact successful. Any mismatch at this time triggers a "Fatal DOS error message. Unable to load COMMAND.COM system halted" and the computer goes to a hard disk dump (FMP 5). This message usually indicates that the path set by COMSPEC is not valid for the message, or also be triggered by differences between the copy of COMMAND.COM used through COMSPEC and the copy from which the checksum was calculated at boot time.

Such differences are most often caused by having installed versions of COMMAND.COM on the system that's having version 5 on the hard drive, but version 5 on the floppy is an even bigger system was worried. These problems may also be caused by patches applied to the disk copy. The patch causes this message if the reboot occurred in the same "clean" stage of the copying of a copy of the checksum, which takes around the patch affects to be calculated.

Another possible cause of the reload error, not even hinted at by the message, is corrupt network situations when the network software redefines COMSPEC to point to the local copy of COMMAND.COM. This copy may differ from the copy that's given to workstation, by causing errors resulting from this problem is to remove the redirection of COMSPEC from the network software, if possible. If redirection is not possible, each workstation must be instructed to perform the network logon. That later, the most fixing COMSPEC to point back to the workstation's own copy.

### The Division Points

When first loaded by the initialization code, up to DOS, COMMAND.COM divides itself into three parts. One part stays where it initially loaded, in low memory, just below the screen, and other programs run. Another part moves to the highest available memory, to be with 100% RAM. The middle portion is discarded after it finishes setting things up. Official DOS manuals distributed with early OI M versions of the system told us that much, but very little else. Here's more of the story.

**Resident, Initial, and Transient Portions** The three parts into which COMMAND.COM divides itself are known as the resident, transient, and initial portions.

The resident portion of COMMAND.COM contains the interrupt service routine, or ISR, that, while it is real-time, is comprised of the interrupt of a separate ISR. In some older versions of DOS it also contains some of the INT 21h service routines, including some of the hardware I/O functions. In addition to these interrupt handlers, the resident portion contains the code that's terminated process termination control using the INT 23 pointer in the PNP. This portion also contains

some initialization data storage associated with the command interpreter, such as the pointers to the transient portion.

The transient portion is loaded only for internal commands. For all others it is freed up for use by a user's command. It is available by space. This area contains the actual input buffer, the code that interprets commands, the internal command list, and all of the code for executing commands. In DOS 5.0 and earlier, if DOS is loaded in the HIMEM.COMMAND.COM, automatically puts much of the transient portion into a HIMEM file rather than at the top of conventional RAM.

Because the actual DOS input buffer is in the transient portion, the standard input editing commands, such as `^Z`, sometimes lose their data if the transient portion is overlaid during the processing of an external command. We'll do another small side trip to look inside DOS Function 05h and see when it happens, where we stop to study the environment. Right now, though, let's continue with the three portions of COMMAND.COM.

Next is the initialization portion. Each time COMMAND.COM is loaded, at least some of this code is executed to verify the file version, load a COMMAND.COM and the resident DOS kernel, search for environment variables, parse any arguments passed to the program. Other actions depend on whether the `^I` option was set, one of the arguments passed. If this switch is absent, the actions described in the next two paragraphs are skipped.

If, however, the switch is present, which is when the primary shell is being installed by DOS, a routine is called that, first, sets the starting address for the transient portion, moves the transient portion into place at the top of RAM, initializes and stores its checksum for future use in reloading, and sets up the interrupt vector for INT 21h to point to the ISR inside the resident portion. Note that this routine also makes this copy of COMMAND.COM the primary shell even if another copy already exists.

In addition, from code there checks for the existence of a file named `U:\D\FNC.BAT` in the current working directory. If it exists, the initialization portion sets flags that direct the input routine to process this file before looking for keyboard input.

In addition, the COMMAND.COM initialization RAM allocation is just adequate for covering up to 16,384 bytes. The copy of COMMAND.COM is not the primary shell, that is the one to which INT 21h points, the resident portion, refers to this copy and not to the primary copy itself. COMMAND.COM then makes control to the command input prompting routine. This routine issues the prompt, waits for input, and dispatches it. But the initialization portion, as such, ceased to exist when the RAM allocation was reduced.

**Where These Portions Are Loaded** The preceding descriptions show you how COMMAND.COM gets itself up for action. In the first edition of this book, I provided detailed instructions on how to create a copy of the way the various parts wind up at your own system. As a result, you can use DOS 5.0 because there are just too many possible variations, depending on what version of DOS you're looking at, on how whether you have UMB enabled, and what model of CPU you're using.

For instance, CompuComp, a major reseller, I found is the author of a forthcoming book on DOS, *Microsoft DOS 5.0 User's Guide*. COMMAND.COM contains a portion of shareable code, extracted from the source files, depending on the model processor and which may have been transferred to the HIMEM file. This address is stored in a calling point table; you use INT 21h Function 5500h to get the table address.

```
INT 21h  AX = 5500h
returns  AX = 0000h
        DS:SI = address of calling point table
```

The procedures use INT 21h's table include services provided for the transient portion, as well as interrupt handlers for INT 7h, 74h, 23h, 71h. For the moment, the procedures manipulate data in the



**COMMAND.COM resident portion** Because of this and the fact that the method by which the dispatcher communicates its data segment is susceptible to change, these procedures must be of only limited use outside of COMMAND.COM itself.

In the first releases of DOS 5.0, they were called from a dispatcher in COMMAND.COM resident portion. The dispatcher was unusually non-reentrant. Its structure is something like this:

```
int_23h:
    push    ax
    mov     ax,0
    jmp     get_address

int_24h:
    push    ax
    mov     ax,1
    jmp     get_address
    .
    .

get_address:
    stl
    mov     cs:ebx_save3,bx      << DANGER !!!
```

BX was then used to index into the table of calling addresses. This opened a window for disaster each time COMMAND.COM received an interrupt because subsequent passage through the dispatcher could overwrite the saved value of BX, which the first interrupt still needed. This would be serious. For INt\_24h were followed a few microseconds later by an INt\_21h or when an INt\_21h interrupted another.

Later copies of DOS 5.0, when disassembled, show that this technique was replaced by individual jumps for each ISR. The table still exists, but it contains far more jumps that can take control directly into the TIMA window for disaster or into . Access is direct, rather than by means of an indexing technique.

### Why Does F3 Sometimes Quit Working?

A few pages back I promised another side trip to discover why the standard input editing keys such as F3 occasionally did a predictable job of working after a large external program set in. The editing function was a fact: all parts of DOS use it, whether set on BAh, whether open, if it then occurs your fancy is intimately connected with the way that COMMAND.COM partitions and reloads itself so that it too much of a digression to examine in detail.

As you have seen, COMMAND.COM spends much of its time parked inside INt\_21\_0Ah. In order for text editing operations to work, Function 0Ah must first get all keyboard input into a temporary buffer. The function can copy the input to the user's own buffer only after input is complete. This temporary buffer is saved on the user's stack, so it vanishes when Function 0Ah returns. To be able to use F3 to reselect input, it is just as if there must be another copy of that line in another copy of a buffer mentioned by Function 0Ah.

It is tempting to jump to the conclusion that this previous copy kept in the DOS data area, but that's not the case. Each application that supports the F3 editing capability must keep its own copy of "last input data" so that input for various applications would be moved in. For instance, you can go into DEBUG by typing "debug" on the DOS command line. In DEBUG you can use F3 to recall characters you type there, but when you return to the command shell prompt and press F3, the last DOS command line "debug" shows up. It has been saved by COMMAND.COM precisely all the time you were using DEBUG.

Incidentally, the official documentation for Function 0Ah is strangely silent so far from the DOS editing functions are brought into play. In fact, the *MS-DOS Programmer's Reference* DOS

• Of course, Function 0Ah "obsolete" and recommends using Function 3h instead, although the Read File Function itself has no editing capabilities at all. It appears, though I have been unable to confirm this yet, that the equipment available to me that the Read File function detects reading from STDIN as a special case and really uses the Function 0Ah code.

The logical requirement for using Function 0Ah is that DOS point to a buffer large enough to accept the expected input starting at its third byte. The first byte of the buffer specifies the total number of bytes it can hold, and the second byte the number of bytes now in the buffer. Normally most programs set the second byte to zero when setting up the buffer.

It is worth noting that this buffer itself serves dual duty as the "raw copy" and as the destination for new input. Several requirements must be met in order to enable function key editing with Function 0Ah, the small assembly language program in Figure 10-5. F3TEST.ASM shows the result.

**Figure 10-5. F3TEST.ASM**

```

title F3TEST
.model small                ; use simplified seg style

;data
buf1    db 126,0            ; buffer for input
        db 126 dup (0)
buf2    db 'This is Buf1 initial string'
buf2l   equ $ - buf2      ; calculated data length
crlf    db 13,10,'$'
buf3    db 128 dup (0)    ; buffer for output

;stack

;code
start   proc
        mov     ax,DGROUP          ; set up segment regs
        mov     ds,ax
        mov     es,ax
        mov     di,offset DGROUP buf1 ; prepare to preload buffer
        mov     si,offset DGROUP buf2

        mov     cx,buf2l           ; get the move count
        mov     al,126            ; establish max size
        mov     ah,cl             ; and current size (essential)
        stc
        inc     cx                ; include the CR in data moved
        rmp     movsb

LT:     mov     di,'!'            ; prompt character
        mov     ah,2
        int     21h
        mov     dx,offset DGROUP buf1 ; get buffered input line
        mov     ah,0Ah
        int     21h
        mov     al,buf1+1         ; check returned length
        cmp     al,1
        jb     L2                ; quit if zero

        mov     si,offset DGROUP buf1+2 ; copy to Buf3 for output
        mov     di,offset DGROUP buf3 ; (essential that BUF1
        xor     cx,cx             ; not be changed in any
        mov     cl,al             ; way, to allow edits)
        rep     movsb
        mov     si,offset DGROUP crlf ; append CR/LF/$
        mov     cl,1
        rep     movsb
        mov     dx,offset DGROUP:crlf ; print CR/LF first
        mov     ah,9

```

```

int     21h
mov     dx,offset BGROUP,buf3    ; then print BUF3
mov     ah,9
int     21h
jmp     L1                        ; go back for another

L2:     mov     ax,4c00h          ; terminate process
int     21h

start   endp
       end   start

```

When you assemble and link this program, then run it, you can press F3 as the end of response to the prompt — just as the preloaded data string that was moved into the input buffer by the code immediately ahead of label L1.

The essential point here is that the length byte at BUF1+4 must point to an ASCII CR character in order to enable the function-key editing feature. Also, the input buffer must either be preserved without change or be restored to the same condition in which function BEdit moved it in order to permit subsequent editing. That's why F3+ESC copies the input line over to buffer BUF3 before appending the CR and LF to it inside the input loop. It's also why the two bytes at BUF1 and BUF1+1 are always added inside the loop but are left just as they were returned by the previous call to F+int on BVA. Simple enough, but it took a while to discover that these were essential to enable the edit capability.

Here's a short sample to show how it works.

```

D:\DOS52\CHAPTER10>F3test
This is Buff initial string
This is Buff initial string
NEW---a Buff initial string
NEW---a Buff initial string

```

```

D:\DOS52\CHAPTER10>

```

On the second row, I simply pressed F3 to echo out the preloaded initial string. I then pressed ENTER (F3+ESC) to reloaded the buffer content back and then prompted me again. This time I typed "NEW" and then pressed F3 followed by ENTER.

The reason why the editing function worked when the transient portion of COMMAND.COM was overlaid by a large program is simple: it's the last input buffer kept by COMMAND.COM internally, the working buffer that moves for parsing the command file. It is located in that transient portion. The primary purpose of this buffer over a disk image of the transient portion is to store what you press F3 or one of the other editing keys after the transient portion has been reloaded from disk. You get only the empty buffer, instead of the actual last input data.

## DOSKEY

The DOSKEY 150K file that comes with DOS provides a command history capability that allows you to recall and re-use previously entered command lines, edit them, and supply them to COMMAND.COM to place a prompt from the keyboard.

As you saw at the start of this chapter when tracing one cycle of COMMAND.COM's operating loop, the shell first checks DOSKEY by means of INT 21h to determine if it's installed and, at the same time, to set a flag to report if DOSKEY is active. Only if DOSKEY is not active does COMMAND.COM call INT 21/30Ah.

DOSKEY maintains a ring-linked collection of input buffers. Each time that DOSKEY is called, it takes the next buffer in the ring, clears its character count to zero, and calls INT 21/08h to get input into that buffer.

The `page` down arrow keys select either a previous or later buffer from the ring, without clearing `ERROR`. This allows your access to any buffer in the collection by repeated use of either arrow key, except, of course, when it's already stepped when stepping in either direction.

When you type to input characters, `DOSKEY` issues the same `INT 28h` as `INT 21h`. All other would.

## Using the Environment

This section discusses the environment as implemented by MS-DOS and used by `COMMAND.COM`. Although the environment itself is simply based at the operating system level (the `INT` contains an environment segment field, for instance), the details of its usage are left to the user and interpreter, with the result that this discussion may not apply fully if an alternate interpreter is used.

**How `COMMAND.COM` Uses the Environment** The idea of an environment for each process had to exist in P.M. (and also exists to MS-DOS from UNIX) but was greatly simplified in the transfer. As a result, as represented in MS-DOS, the environment consists of a paragraph-aligned block of space that is roughly `32 * 67` bytes in size. In addition, as a practice, it is always much smaller. This block is made up of a collection of environment variables, each of which consists of a variable name followed by the value of the data, both in a name and the data are stored as an `ASCII` text string, with an equal sign separating the name from the value. Note that at this is just the convention used by `COMMAND.COM` and not something built into `DOS` itself.

`DOS` will allow a block to be used by a process and its environment. `DOS` puts the segment address of the environment into the `environment` field of the `INT` when `DOS` dispatches the process through `DOS EXEC` `INT 21 4bh`.

The program that calls `EXEC` is responsible for allocating the environment space and defining all the variables to be used. It is up to the user to define. If the segment address passed to the `DOS EXEC` function is `0`, then `DOS` will allocate as much space as it can to contain a copy of each variable in the current environment and store it in the system automatically. The size used (unformatted) except for the overhead of the environment space used internally by the primary shell must be set up explicitly by the program shell itself without variables. Otherwise, no master copy of the environment will exist.

`COMMAND.COM` uses several predefined environment variables to control its actions. `COMMAND.COM` uses `TIME`, `DATE`, and `TIME` specification that is used each time `COMMAND.COM` must re-allocate its environment. `PATH` lists the names and paths to be searched for possible external commands. `COMP` is the string of characters that `COMMAND.COM` uses to prompt for user input.

Both `TIME` and `DATE` use separate internal commands that can be used to modify their contents. However, the environment variables can be modified by the internal command `SET` if the variable does not exist or is not set, with the environment provided that enough space exists for it. From the primary environment space, the `SET` command operates on the master environment.

The master environment variables `TIME`, `DATE`, `PATH`, `COMP` only are created in the master environment if the program uses the primary shell with internal code. The size of the master environment is also determined at this time. It is 160 bytes by default, but can be altered by the environment in the `SET` field. `COMMAND.COM` uses `CONSOLE` `SYN`. Alternate shells use different syntax but have the same capability.

Once the primary shell has been loaded, the environment space allocation cannot be increased. While it might appear simple to replace it with a new larger copy and adjust the segment address at

PSP 002Ch record why that original master environment address was also saved in unimplemented internal locations for use by internal commands. Those addresses would be difficult to locate & labels in order to modify them.

When COMMAND.COM dispatches an external command, it simply passes the file name to DOS, thus generating an exact copy of the master environment for use by the spawned process. Because it's a copy and not the original, any changes made to it by the process will not be reflected in the master environment itself. Although this protects the master environment from being altered accidentally, it's still a bit of an annoyance, except by using the SETC command from the primary shell command-line prompt, which is not always convenient.

**Locating the Environment** Before you can make a change to the information stored in the environment area, you must first find it. DOS stores the segment address of the environment area for each process in the word at offset 2Ch in the PSP, but if you need access to the master environment, you must locate the PSP for the primary shell.

The assembly language package shown in Figure 10-6, ENVPKG.ASM, contains three routines designed to support small model C programs. Curenvpt and mstenvp locate the current and the master environment areas, respectively, and envsiz locates the first free address. envsiz is input a pointer, such as the one returned by the first two, and returns the size of the area in paragraphs.

**Figure 10-6. ENVPKG.ASM**

```

;ENVPKG.ASM  _In Kyle - July 1990, April 1993

.model small,c

.data
; assumes being used from C with _psp global variable
    extrn  _psp word

.code

curenvp proc
public curenvp
; char far * curenvp( void );
    mov     ax,_psp           ; get PSP seg
    mov     es,ax
    mov     dx,es:[002Ch]    ; get env address
    xor     dx,dx           ; offset is zero
    ret
curenvp endp

mstenvp proc
public mstenvp
; char far * mstenvp( void );
; This is the only guaranteed method for COMMAND.COM if DOS=HIGH is
; in effect, but has problems under some circumstances.
; See notes in text
    mov     ax,352Eh        ; get INT2E vector
    int     21h            ; vector in ES:BX
    mov     dx,es:[002Ch]   ; es:2C is seg of environment
    xor     dx,dx          ; make the offset zero
    ret
mstenvp endp

envsiz proc    oenv:word, senv:word
public envsiz
; short envsiz( char far * vptr),
    mov     ax,senv        ; get segment of env
    dec     ax             ; back up to RCB

```

```

mov     es,ax
mov     ax,es [0003h] , get size in paragraphs
ret
envsiz  endp
end

```

The `envsiz` routines following each of the `publi` directives are simple prototype declarations to be used by the linker program that uses `ENVPKG`. To use these routines, first assemble the program into an `OBJ` file as follows:

```

MASM /ms ENVPKG,
TAGN /ms /j/RASPS1 ENVPKG,

```

The `envsiz` routine used by both assemblers forces procedure names to remain in lower case so that the `OBJ` file can be linked into non-`dos` programs. The operation of `envsiz` and `envsizep` also relies on the fact that compilers for the PC return four byte far pointers in the `DX:AX` register pair.

The `envsizep` routine assumes an external global variable called `psp`. This variable is provided in `dllc` environments for the PC using Borland compilers, `psp` in `DOS 3.1` and others in `NTDIB.H`. You could change this to use `DWORD` or `INT32` if you prefer to avoid compiler differences.

The `envsiz` routine is based on the fact that every environment block starts at an offset of zero and contains a 16-bit Memory Control block that contains the size of the block in paragraphs. See Chapter 11. Thus, when you pass a far pointer to the environment block to `envsiz`, it increments the segment of the address by its own `MBR` and then retrieves the size from offset 3 in that segment. Note that the returned value is always a paragraph. If a byte size is needed, the returned value must be multiplied by 16 (shifted left 4 places).

The sample program `EPTST.C` shown in Figure 10-7 uses the `ENVPKG` routines. Compile this with the standard memory model and link with `envpkg.obj`.

**Figure 10-7 EPTST.C**

```

;*****
; EPTST.C Tests environment-access object modules
; Jim Kyte, July 1, 1990, April, 1993
;*****
#include <stdio.h>

char far * curenvp( void ), /* prototype declarations */
char far * mstenvp( void );
short envsiz( char far * vptr);

void main( void)
{ char far *mine,
  char far *master

  puts( "\nEnvironment locations :");
  mine = curenvp();
  master = mstenvp();
  printf( "Current environment is at %p, size- %i bytes\n",
  mine, envsiz(mine)<<4 );
  printf( "Master environment is at %p, size- %i bytes\n",
  master, envsiz(master)<<4 );
}

```

Typical output from `EPTST` when run at the primary shell's command prompt, follows:

```

Environment locations:
Current environment is at 1868:0000, size: 256 bytes
Master environment is at 1107:0000, size: 512 bytes

```

This shows how the current working copy of the environment has been trimmed to a size just adequate to hold the defined variable strings and the program's pathspec.

**Other Ways of Locating the Environment** The preceding discussion has a problem which does not often arise but is serious when it does. The interrupt vector for INI 21h does not always point to the primary copy of the command interpreter. For example, using David Macey's INTRMP program from Chapter 5, it is trivial to take over INI 21h for the purpose of seeing where `system() & instcomp` comes along while INTRMP is in control of INI 21h. It would think that INTRMP is the primary command interpreter, even though INTRMP has simply hooked INI 21h for diagnostic purposes and is just changing the interrupt to the previous owner, presumably, the primary command interpreter.

Of course, this is unlikely to happen out in the field. After a 16-bit morning, but a command interpreter really takes over INI 21h, right? Still, it's easy to fool `instcomp`, so it is difficult to know, or know how to do so.

What other ways are there of locating the master environment?

One technique avoids the issue of locating the master environment entirely. It issues M1 commands using INI 21h. The ability to call COMMAND.COM using INI 21h is discussed later in this chapter. Note that `instcomp` does not call INI 21h at which issue its interrupt vector, so hopes of finding the primary command interpreter. This technique is repeated in `instcomp` in MSINSP.C below.

Another, different technique, used by the BVT compiler, discussed earlier in this chapter, walks the MCB chain looking for the environment segment belonging to the primary command interpreter. Chapter 7 documented how to walk the MCB chain; the key here is to take the first environment you find in the chain.

How do you find the first environment? Don't look for certain AMB characters to see if an MCB corresponds to an environment; instead, look for PNP; walk the PNP chain, as it were, and look at offset 2Ch. Take the first environment you find. How do you know you've got a PNP? Well, look for the escape code bytes for INI 20h, 4Dh, 20h, like some programmers, instead, look for an MCB whose offset is greater than the MCB offset. See the routine `system2` in Figure 10-8.

Another technique, designed especially to accommodate command interpreters loaded with the program, also walks the MCB chain. This time, you look for the key character `0x01010101` rather than its corresponding PNP. For a moment, why? You look for the first environment because of the program and you look for a character not higher than its PNP because the command interpreter loads an environment for itself. Therefore, the environment is at a higher address than the program. See `system2` in Figure 10-8. Note that this technique can cause trouble if you load programs into UMBs, since some of the memory manager routines that it supplies the environment affect the program rather than force `system` to load high. It also fails if you use 386s and have it swap itself into a UMB.

Finally, another technique involves walking back along the PNP chain until you find PSE that is its own parent. However, this technique, which first appeared as Jerry Stroup's "Providing Program Access to the Real DOS Environment," *PC Magazine*, 28 November 1989, pp. 309-314, is designed to find what we called the active environment, not the master environment. This method involves the currently active shell. It works with a program spawned from the statement in a real-mode batch program; for example, you would find the master environment if using a technique.

On the other hand, this is the only reliable way to locate the effective active environment if you are running in a DOS box from Windows, since each DOS box gets its very own copy of the master environment. For example, ENVDEL, as presented in our first edition, did not work in a DOS box under Windows. It affected the wrong environment, one too high up in the chain. We've replaced ENVDEL to overcome this problem in the current version.

Three techniques for finding the master environment appear in the routine in Figure 10-8. The function walk() in MSTENV.C is a variation on the UDMEM.C program from Chapter 7. Here, however, walk() expects a pointer to a function. For each MCB it finds, walk() calls this function, passing it the MCB pointer. The function should return TRUE to indicate that walk() should keep going, or FALSE, to indicate that walk() should stop.

Thus, you can plug different functions into walk() making MSTENV.C (Figure 10-8) a test bed to compare and differentiate methods for locating the master environment using the MCB chain.

**Figure 10-8. MSTENV.C**

```

/*
 * MSTENV.C
 * Test bed for different methods to find master environment
 * Andrew Schulman, July 1990
 * modified July 1992 to use undocdos.h typedefs
 * and also for UMB support - Jim Kyle
 */

#include <std.lib.h>
#include <dos.h>
#include "undocdos.h" /* see Chapter 7 */

extern LPMCB firstcb(void);

,*****
LPMCB (ptr mcb;void) /* returns far pointer to first MCB */
{
    ASM mov ah, 52h;
    ASM int 21h;
    ASM mov dx, es:[bx-2];
    ASM xor ax, ax;
} /* in both Microsoft C and Turbo C, far* returned in DX:AX */

char far *env(LPMCB mcb) /* returns far ptr to env, or NULL */
{
    char far *e;
    unsigned env_mcb;
    unsigned env_owner;

    if (! IS_PSP(mcb))
        return (char far *) 0;

    e = RC_FP(ENV_IR_PSP(mcb->owner), 0);
    env_mcb = MCB_FP_SEG(FP_SEG(e));
    env_owner = ((RCB far *) RC_FP(env_mcb, 0))>owner;
    return e;
}

typedef BOOL (*WALKFUNC)(LPMCB mcb);

/*
 * General purpose MCB walker
 * The second parameter to walk() is a function that expects an
 * MCB pointer, and that returns TRUE to indicate that walk()
 * should keep going, and FALSE to indicate that walk() should
 * stop.
 */
BOOL walk(LPMCB mcb, WALKFUNC walker)
{
    for (;;)
        switch (mcb->type)
        {
            case 'M':
                if (! walker(mcb))
                    return FALSE;
                mcb = (LPMCB)RC_FP(FP_SEG(mcb) + mcb->size + 1, 0);
                break;
            case 'Z': /* end of an MCB chain */

```



```

        walker(mcb);
        if( FP_SEG( mcb ) < 0xA000 )
            ( mcb = firsth(), /* bridge to UMB chain */
              if( mcb == ( FP ) 0L ) /* none found */
                return FALSE;
            )
        else
            return FALSE; /* UMBs all done too */
        break;
    default:
        return FALSE; /* error in MCB chain */
}

)

/*****
/* using the GETVECT 2Eh technique (ENVPRG ASM) */
void far *wstenvp1(void)
{
    ASM mov ax, 352eh /* get INT 2Eh vector */
    ASM int 21h
    ASM mov dx, es, 1002ch /* environment segment */
    ASM xor ax, ax /* return far ptr in DX:AX */
}

/*****
/* walk MCB chain, looking for very first environment */
void far *env2 = (void far *) 0,

BOOL walk2(LPRCB mcb)
{
    env2 = env(mcb);
    return env2 ? FALSE : TRUE;
}

void far *wstenvp2(void)
{
    walk(get_mcb(), walk2);
    return env2,
}

/*****
/* walk MCB chain, looking for very LAST env addr > PSP addr */
void far *env3 = (void far *) 0;

#define NORMALIZE(fp) (FP_SEG(fp) + (FP_OFF(fp) >> 4))

BOOL walk3(LPRCB mcb)
{
    void far *fp,
    /* if env seg at greater address than PSP, then
       candidate for master env -- we'll take last */
    fp = env(mcb),
    if (fp)
        if (NORMALIZE(fp) > (FP_SEG(mcb)+1))
            env3 = fp;
    return TRUE, /* always return TRUE to get last */
}

void far *wstenvp3(void)
{
    walk(get_mcb(), walk3);
    return env3;
}

/*****
/* walk MCB chain, looking for very first environment belonging
   to PSP which is its own parent */
void far *env4 = (void far *) 0,

BOOL walk4(LPRCB mcb)

```

```

{ env4 = env(mcb);
  if (env4)
  { unsigned psp = FP_SEG(mcb) + 1;
    return (PARENT( psp ) == psp) ? FALSE /*found it!*/ : TRUE;
  }
  else
    return TRUE; /* keep going */
}

void far *mstenvp4(void)
{ walk(get_mcb(), walk4);
  return env4;
}

/*****
void main( void )
{ printf("GETVECT 2eh method; mstenvp1 = %fp\n", mstenvp1());
  printf("WALK RCB method; mstenvp2 = %fp\n", mstenvp2());
  printf("WALK RCB/LAST method; mstenvp3 = %fp\n", mstenvp3());
  printf("WALK RCB/OWN PARENT; mstenvp4 = %fp\n", mstenvp4());
}
*****/

```

How does this program behave under various conditions. First let's try it from the normal COMMAND.COM prompt:

```

D:\DOS2\CHAP10>mstenvp
GETVECT 2eh method; mstenvp1 = 1A7B 0000
WALK RCB method; mstenvp2 = 1A7B 0000
WALK RCB/LAST method; mstenvp3 = 1A7B 0000
WALK RCB/OWN PARENT; mstenvp4 = 1A7B 0000

```

Not so good. All four methods agree on where the master environment is. Let's see what happens if we load a child copy of the command interpreter:

```

D:\DOS2\CHAP10>mstenvp
GETVECT 2eh method; mstenvp1 = 1A7B 0000
WALK RCB method; mstenvp2 = 1A7B 0000
WALK RCB/LAST method; mstenvp3 = 1B29 0000
WALK RCB/OWN PARENT; mstenvp4 = 1A7B 0000

```

Oh-oh! The third method picked up the child copy's environment instead of the master's. But that might not always be bad because sometimes you need to be able to locate the active rather than the master copy.

Now let's load a new primary command interpreter with the `p` flag and run MSTENVP again:

```

D:\DOS2\CHAP10>mstenvp
GETVECT 2eh method; mstenvp1 = 1B29 0000
WALK RCB method; mstenvp2 = 1A7B 0000
WALK RCB/LAST method; mstenvp3 = 1B29 0000
WALK RCB/OWN PARENT; mstenvp4 = 1A7B 0000

```

Here `mstenvp2` and `mstenvp4` were correct. They both stack with the old abandoned environment segment 01A7B0000. Instead of upgrading as the other two subfunctions did. The `p` flag on `COMMAND.COM` or `USER.COM` or `NLS.COM` creates a new primary command interpreter and a secondary command interpreter. Strika out for `mstenvp2` and `mstenvp4`.

What happens if you load a second child copy of the command interpreter under the new primary? By now you should be expecting to find at least three different environments and you do:

```

D:\DOS2\CHAP10>mstenvp
GETVECT 2eh method; mstenvp1 = 1B29 0000
WALK RCB method; mstenvp2 = 1A7B 0000
WALK RCB/LAST method; mstenvp3 = 1B07 0000
WALK RCB/OWN PARENT; mstenvp4 = 1A7B 0000

```

Only `mstenvp4` correctly located the real master environment; in this case, `mstenvp3` grabbed the child copy just as it started, while the other two techniques were still stuck on the original, no longer valid, version.

After loading a debugger (INSTRN), with an appropriate script that hooks INI 21b, we get this:

```
D:\DOS2\CHAP10>mstenvp
GETVECT 2Eh method, mstenvp1 = F800 0000
WALK_MCB method, mstenvp2 = 1A7B 0000
WALK_MCB/LAST method, mstenvp3 = 1B29 0000
WALK_MCB/OWN_PARENT, mstenvp4 = 1A7B 0000
```

Now, `mstenvp1` is wrong, i.e., Segment F800 points into the debugger, not the primary copy and interpreter. Only `mstenvp2` and `mstenvp4` properly identified the master environment, and they too would have been wrong had a second primary been present.

Finally, reboot the machine, load a TSR (i.e., CONERG.MSY with the INI.MFI command) and try `MSTENVP` again:

```
D:\DOS2\CHAP10>mstenvp
GETVECT 2Eh method, mstenvp1 = 1C35 0000
WALK_MCB method, mstenvp2 = 1968 0000
WALK_MCB/LAST method, mstenvp3 = 1C35 0000
WALK_MCB/OWN_PARENT, mstenvp4 = 1C35 0000
```

In this situation, `mstenvp2` was wrong again. Segment 1968 points into the cyrcorant of the TSR code I used with the INI.MFI command. This is any program that uses the "walk\_MCB" method of finding the master environment (only whenever a user of DOS 4.0 or higher uses the handy INI.MFI command).

Could a user's environment be made to "mstenvp2" to detect this situation? Yes. Take out the first environment you find, but the first environment belongs to a TSR that's its own parent (another "mother"). Let's let `mstenvp4` option show all the sample runs.

This works better but only for DOS 4.0 and higher. INI.MFI started out as a command rather than a utility and was the product of a "prototype" that caused `mstenvp2` to fail. It also has miserable quality. Quality 386MMX is very suspect, which some users after the original copy of MCLive Upper Memory Blocks exchanges makes are defined in the UMEMEM file of chapter 7. Even a master copy of 386 appears to be "stacked" (i.e., under 386MMX) it won't ever be found by this technique!

All these methods fail if someone has used a device driver to change the memory allocation strategy to not include the top-level stack loaded. In such a case, you may have to settle for finding the active copy, or, rather, the real master copy, by working back up the parent process chain, as shown in a ROBUST program of Chapter 7 or in the section that follows.

They, although all these techniques work pretty well in many cases, none presents a 100 percent fool proof method for finding the master environment. Of course, the problems I saw, others are the result of some hardware conflict or special cases, but trying to anticipate system cases without doingishes, a professional programmer from the outset. So what's a programmer to do? One technique, of course, is to perform the `mstenvp` function in two or more different ways, compare the results, and bail out if they don't match.

If you're stuck, you'll find that Microsoft will mess up finding the master environment if it's own APPEND file, the APPEND.E switch, was executed in a second, or a command interpreter, does not affect the master environment. In fact, in fact, some system cases (i.e., DOS 4.0), these crashes may also be due to APPEND's interception of FCB open and get file size requests which corrupt register DV. (The file is not actually found, a bug which Control C happens to be present since at least DOS 3.30).



not most implementations make still another UNIX-style copy of the variables into their own data areas. This copy supports the optional copy argument to `mapn`.

For maximum flexibility, you need to be able to access any copy you desire. So when you happen to be shielded out of a program, changes made to the current copy vanish without a trace when you return to the parent program. Changes made to the master copy remain, but they will have no effect until you return to the primary shell.

Several methods could be used to provide totally flexible access to the environment. In one, we'll explore, is designed for easy expansion to other needs. Its foundation is the small assembly language function shown in `NXTEVAR.ASM` (Figure 10-9). This function accepts as input a pointer to one environment variable; it returns either a pointer to the first byte of the following variable (which may be all zero bytes) or `NULL` if the input pointer already points to an all-zero byte.

**Figure 10-9. NXTEVAR.ASM**

```

; NXTEVAR.ASM - Jim Kyle - July 1990

.model small,c
.code

nxtavar proc    uses di, vptr:far ptr byte

    public nxtavar
; char far * nxtavar( char far * vptr ),
    lea    di, vptr
    mov    cx, $000h
    xor    ax, ax           ; search for 0 and...
    mov    dx, ax           ; ...initialize return 0:AR to 0:0
    repne scasb            ; search ES:DI for char 0 in AL
    inc    cx              ; CX = $000h if only one 0 found
    js     nev             ;
    mov    dx, es
    mov    ax, di
nev:    ret
nxtavar endp
end

```

This function, a script prefixed with a far pointer to any ASCII string (not just one in the current area), returns a far pointer to the byte that follows the end of string (not just byte). In the environment's structure, that pointer is either to the first byte of the next string or to a byte that is all zeros. In the latter case, the end of the environment's variable area has been reached (so another call to `nxtavar`, using that pointer, would return `NULL`).

Before that happens, the previous call to `nxtavar` will have returned the far pointer to the all-zero byte. If each pointer is stored in an array, the final one can be used to remove the process's path name. If a final pointer is also can be passed to a routine that first verifies that the next two bytes are 0FF and 00h (a key's speciality with no argument from the regular files). Because the path data is a C string, analysis is a ASCII string; the normal C string functions deal with it properly.

This function is flexible because you can pass the function a pointer to any environment field and the current environment is obtained with the concept of a master of `ENVPKT`, the master environment you obtain by using `strncpy` of the active environment which I used in chapter 5 to obtain with the `ENVEDIT.C` program.

To use `NXTEVAR.ASM`, you must assemble it into an OBJ file following the procedures noted earlier in this chapter for `ENVPKT.ASM`. To illustrate how `NXTEVAR.ASM` is used, Figure 10-10 shows a little C program that reports the location and contents of each string in the current environment.

Figure 10-10. NEV.C

```

/*****
 * NEV.C - Next Environment Variable
 * Jim Kyle, July 1, 1990
 *****/
#include <stdio.h>
#include <stdlib.h>

char far * nstovar( char far * vptr );
char far * curenv( void );

void main (void)
{ char far * menv,
  * envp = curenv();
  while ( menv )
    { printf("Env Var at %p: %s\n", menv, menv );
      menv = nstovar( menv );
    }
  exit(0);
}

```

To compile NEV.C to a C program, use either Microsoft or Borland C compilers, both the `ENVPRG.OBJ` and `NATIVE.OBJ` modules must be used this time because NEV.C calls `curenv()` to locate the current position of NEV.C calls you the address and contents of each variable string in the current system in the sequence in which they occur in the environment block. More often, you will use `getenv()` to specify environment variable by name, and you may want to locate it in the master environment rather than in the current copy. You could use the Library function `getenv()` to return a pointer to the named variable in the current environment copy, but `getenv()` won't access the master.

To both search by name and use the master environment, we modify NEV.C into a far more useful program, FMEV.C, as shown in Figure 10-11. This also requires `ENVPRG` and `NATIVE`.

Figure 10-11. FMEV.C

```

/*****
 * FMEV.C - Find Master Environment Variable
 * Jim Kyle, July 7, 1990
 *****/
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

char far * nstovar( char far * vptr ),
char far * mstenv( void );

void main ( int argc, char * argv[] )
{ char far * menv,
  char vname[128], *vdata, tgt[64],
  int tlen;

  menv = mstenv();
  if ( argc < 2 )
    { printf("Var to find: ",
      gets( tgt );
    }
  else
    strcpy( tgt, argv[1] ),
    tlen = strlen( tgt );

  while ( menv )

```

```

( sprintf(vname, "%s", menu );
  if ( vname[0] == 'a' )
    ( vdata = &name[0],
      vname[0] = '\0',
      if ( strcmp( tgt, vname ) == 0 )
        break;
    )
  menu = nextenv( menu );
)
)

if ( menu )
( printf("Found %s at %p\n", vname, menu, vdata ),
  exit(0),
)
else
( printf("%s not found.\n", tgt );
  exit(1);
)
)
)

```

In EMX, the declaration of `main` has been changed to permit the desired variable name to be entered as a command-line argument. Using `sprintf` with the `%s` format specifier forces the library routine to do all necessary conversion to copy each string argument from the environment to the local work area. Although newer compilers include the `strcpy` library function to do some of the pointer-pointing tasks, the `sprintf` solution still seems much simpler to comprehend.

We'll return to the environment at the end of this chapter, with the ENVIRONMENT utility.

### INT 2fh, the Back Door

Interrupt 2fh provides a back door into COMMAND.COM, allowing access to both the main interpreter routine and the command dispatcher from outside the program. Control is passed to the latter function if not used by the command interpreter for its purpose. Otherwise, it is used.

The procedures to which INT 2fh provides access, though, are not arbitrary. The execution of batch files and the FORTRAN.COM file are considered as well as these procedures. Consequently, an external program that calls INT 2fh can't do anything but call it. It's important to note that this interrupt is serviced by code in the resident portion of COMMAND.COM itself rather than by DOS. If an alternative command interpreter such as 4DOS is installed in place of COMMAND.COM, the description provided in this section will not be true.

Although INT 2fh can be dangerous if you use it without knowing its limitations, it does permit access to the primary copy of the command interpreter skill file, one dispatched by CONFGSYS during system bootstrap.

The real meat so to speak aspect of and/or related DOS can be found in Daniel I. Greenberg's article "Recovering the DOS Shell" which appeared in *Pro Systems Journal*, May 1, 1990 on pages 28-36. It also features a definitive piece on INT 2fh. Unfortunately, *PSJ* is no longer published, and copies of this issue may be nearly impossible to locate.

**The Function** The purpose of INT 2fh is to provide external programs with a method for accessing the command parsing and dispatching routines within COMMAND.COM. It is a weakness which copy of COMMAND.COM is in a primary shell if no other copies exist on the system. The copy to which the INT 2fh access points is defined as primary.

The use of INT 2fh is quite simple; programs to use COMMAND.COM's internal routines take SET to modify environment variables. Since it invokes both the interpreter and dispatcher routines of COMMAND.COM, INT 2fh can be used to launch a batch file. However, its use is highly restricted by the fact that COMMAND.COM is not recursive. Any program that uses it to launch a batch file will take a batch file and directly launches another, never returning to its own special

can use only COMMAND.COM supports this interrupt, it's only prudent to verify that the interrupt service routine exists before attempting to use it.

The assembly program HAVE2F.ASM shown in Figure 10-12 provides a C-callable routine that returns IRET if the first byte of the service routine for INT 2Fh is not the IRET that DOS puts in the vector by default; the routine returns FALSE if the byte is equal to the IRET opcode.

Figure 10-12. HAVE2F.ASM

```

model small,c
code
have2e proc                ; returns 1 if ISR exists, else 0
public have2e
int have2eC void ; /* prototype */
mov     ax,552Eh
int     2fh
mov     ax,es
; test for empty just to be safe
or     ax,bi
; although MS-DOS never leaves this
; vector all zeroes
jz     hi
mov     al,es.[bx]
xor     al,0CFh
; opcode for IRET
jz     hi
mov     ax,1
hi:
cld
ret
have2e endp
end
  
```

**Its Use** The INT 2Fh takes a command string such as DIR \*.\* or anything else a user might normally type at the prompt, puts it in the buffer and a byte count at the front, just as INT 21h (IO) does, and then invokes the command. The service routine for INT 2Fh. The command is executed. But to use INT 2Fh successfully, you must follow several rather strict guidelines. This capability has never been officially sanctioned. Microsoft's *MS-DOS 5.0 Programmer's Reference* describes INT 2Fh merely as "a debugging option" for use "in COMMAND.COM only," despite the fact that it's never called within COMMAND.COM. As a result, INT 2Fh wastes most of the system trapping abilities of more normal functions.

The most important restriction on the use of this interrupt is that no registers, not even the stack segment and stack pointer, are preserved. Immediately at entry to the service routine, the return address and a count are popped from a dedicated storage area in COMMAND.COM's own data segment. On completion of the task assigned by using INT 2Fh, control returns to the caller by means of a return to system saved segment register values.

It's also important to remember that, just as always, when you call INT 2Fh, you must save a return address, a location that never changes, prior to an interrupt. CS and IP are valid. Notwithstanding the fact that INT 2Fh is a system interrupt, it's not interruptible. That is, if it's called a second time, the second call destroys the return address for the first invocation, leading to system lockup.

It's possible to overcome the lack of a return address by copying critical data areas from the COMMAND.COM data segment to a dedicated memory space, contained specifically for the purpose, then using INT 2Fh and explicitly restoring the data areas from the saved copy. This is essentially how the program COMMAND.CAL works so that it can operate successfully from within a batch file.

The program DO2F.ASM shown in Figure 10-13 contains the function do2e, which sets everything up properly for invoking the interrupt, executing a command string, and regaining control upon return.



**Figure 10-13. DOZE.ASM**

```

.model small,c
.code
do2e    proc    uses ds si di, cmdstr:ptr byte
        public do2e
; void do2e( char * cmdstr ); /* prototype */
        push   ds                ; save data segment regardless of model
        mov    si,cmdstr        ; small model
;      lds    si,cmdstr        ; large model
        mov    cs:svss,sp       ; save stack pointer
        mov    cs:svsp,sp
        cld
        int    2Eh              ; issue command
        cli
        mov    si,cs:svss       ; restore stack pointer
        mov    sp,cs:svsp
        pop    ds                ; restore data segment
        ret
        even
svss    dw    0
svsp    dw    0
do2e    endp
        end

```

This routine should not be used unless `have2e` returns TRUE to indicate that the interrupt is in fact present in your system. The command string passed to the `do2e` function must be in a special format: the same format in which `C_COMMAND.COM` expects to be given. The first byte of the string must contain a binary count of the string's length (excluding the count byte and the terminating carriage return character), and the string must be terminated by a 00h CR character.

Although you can build your own routines using only the `do2e` function of DOS21.CBI, it's easier if you work through an intermediate-level support routine such as the `C` function shown in Figure 10-14.

**Figure 10-14. SEND2E.C**

```

/*****
 * Send2E.C - support for INT 2Eh
 * Jim Kyle, July 1990
 *****/
#include <stdio.h>
#include <string.h>

int have2e( void );
void do2e( char * cmdstr ); /* prototype */
/* prototype */

int Send2E( char * command )
{ char temp[130];
  int retval;

  if( !have2e() )
    ( sprintf( temp, "\x2e\r", strlen( command ), command );
    do2e( temp );
  )
  return retval;
}

```

This snippet of code takes just the command line itself, as you would type it at the prompt, and adds the carriage return and termination, CR. It then passes the edited string on to IN12H using the assembly language instruction. These actions happen only if the interrupt is present; if not, they are skipped. Finally, Send2E returns TRUE if the command was passed to IN12H and FALSE if it was not.

The program in Figure 10-15 runs Send2E in a loop, making a crude command interpreter.

Figure 10-15. TEST2E.C

```

/*
TEST2E.C
    Revised August 1992 - jk
Turbo C++ 2.0:
    gcc test2e.c send2e.c do2e.asm have2e.asm
Microsoft C 6+:
    cl -gc test2e.c send2e.c -RMax do2e.asm have2e.asm
Borland C++ 3+:
    bcc test2e.c send2e.c do2e.asm have2e.asm
*/

#include <stdio.h>
#include <stdlib.h>
#include <string.h>

main()
{
    char buf[128],
    for (;;)
    {
        fputs("\n", stdout),
        gets(buf),
        if ((strcmp(buf, "bye") == 0 || strcmp(buf, "BYE") == 0) ||
            (strcmp(buf, "exit") == 0 || strcmp(buf, "EXIT") == 0))
            break,
        Send2E(buf),
    }
    puts("Bye"),
}

```

We already saw this program in Chapter 9 on ISRs, where it was used with Ray Michels' ISR skeleton to create a user-resident command interpreter, ISR2E.

It is worth noting that this program, and IN12H more general, is compatible with the installable core mode. For those interested in this subject, for example, if you run our INSMID program, you can specify a command line command FILE NAME using an IN12H program such as IN12E or ISR2E.

IN12C also iterates in default case, going to the command loop of COMMAND.COM, which then iterates with the COMMAND's child program. So, this is only to be expected. It is worth noting that IN12C is subject to the re-entrancy question. IN12C supports a situation one expects to obtain from the command processor.

An interesting thing Chapter 9's program uses Send2E to take a M/D command. The master environment, instead of an overview of the environment, belonging to the program, is updated. Using IN12E is thus one technique for updating the master environment; other, safer, techniques were detailed earlier in this chapter, in the section "Using the Environment."

Notice that the issue of byte order was not addressed in Send2E. That's because the issue is much more complex than just saving the data areas and restoring them. Although that's necessary, other considerations must also be addressed—like IN12H's calls to boot enough for dependable use.

The first question that arises is how much information from the DOS area. Greenberg's article in *Programmer's Journal* referenced at the start of this section suggested that 120 bytes would include all necessary information for DOS versions 2.0 through 4.01. However, not all of that 120 bytes

needed to be restored, and the locations that required restoration varied from one DOS version to the next. DOS 5.0 and 6.0 had not yet appeared when Greenberg's article was published, but the 120-byte figure can be used as a starting point for experimentation.

A major question equally important concerns the methods required to handle CTRL-BREAK interruptions or a local error condition. In both cases, the normal COMMAND.COM response is inadequate to provide full safety.

Since the first edition of this book was published, the issue of running INI files (such as the HIMEM.EXE) from within a batch file has garnered some publicity. Even before that, Michael Methane's "Running Programs From a Batch File" (*PC Magazine*, February 16, 1988) had written that programs using INI files "will not execute batch files if a crash occurs from within a batch file."

Jeff Proulx, in a 2004 article on undocumented DOS 5.0, also indicated DOS 4 and 5 only "if you attempt to crash your system with every combination situation. The reason one of the programs you are using is crashing is due to a batch file. Since INI files don't execute if DOS uses a certain batch file, so if you run a batch file as a INI file from your program, your system will crash."

We can have had problems running programs that use INI files from within a batch file, but nothing so dramatic as crashing the system. Instead, we may found simply that INI files not loaded properly and that a certain error message. As it was found no indication at all of some disassembly of COMMAND.COM from DOS 5.0 that DOS uses it to run batch files. INI files is not involved at all with a COMMAND.COM. Its entry point, however, jumps to the self-same internal code used for batch file processing.

It would be nice if we could develop this issue of INI files and batch files once and for all and say what files and files not work, and why. However, doing so would take far more effort than the rest of warrants since it appears likely that at least in other version of DOS will be appearing annually. Any result is spent on DOS 5.0 or 6.0, it is not only version-specific, and would probably fail to apply to the next version.

Greenberg presented a major program written in assembly language to do a separate attack, and a second one is recommended to wait to use INI files as a way for security programming. However, because there revisions of DOS will in fact really use the system, this capability is cast as significant as well, but for a complete of each revision of the past, we can say, consist of avoid the use of INI files completely is to be found on, including you that they were of the undocumented internal workings of COMMAND.COM. It is needed to execute external commands from within your program, use functions such as system() or the spawn() family, etc.

## Alternatives to COMMAND.COM

Several alternatives to COMMAND.COM now exist, and this section looks at a sampling of them. Some provide a simple replacement of the command interpreter, while others claim to be allowing control of interpreter behavior, to command-line interface from the user.

### 4DOS.COM

This command interpreter from T.P. Software totally replaces COMMAND.COM. Originally distributed only as shareware, it is now also available through retail outlets. As this was being written, 4DOS.COM was at version 4.02.

Another alternative is NDOB, distributed by Sinterax with recent versions of the Norton Utilities. This is a very slightly modified version of 4DOS version 3.03, which has been customized for interaction with the NU's enhanced capabilities, but which does not have 100 of the features added to 4DOS in its version 4. The NDOB supplied with NU 4.0 is a newer version, fully DOS 6.0 compatible. Most of the description of 4DOS that follows applies to NDOB also.

Other products that replace COMMAND.COM include Command Plus, PolyShell and DYNASHL. I have chosen to describe 4DOS because it typifies the group and because I use it daily and to provide a comparison with the others. But all of the programs offer improved capabilities when compared to COMMAND.COM.

**A Total Replacement, Plus More** Like all of its competitors, 4DOS provides near total compatibility with COMMAND.COM, even going so far as to duplicate strange actions that most users consider "bugs." For example, to minimize the number of applications that can run without change.

Where COMMAND.COM has only 39 internal commands at most, however, 4DOS now provides more than 80 third-party making many utilities used with COMMAND.COM obsolete. For example, this command and its parent includes a built-in environment editor, ENVI, which allows you to edit either the master or the active environment.

For building force-free command systems, internal commands are able to draw boxes on the screen in color, input from the keyboard that then modifies batch file execution decisions. It's even possible to perform arithmetic using the EVAL function and to peg the result into a command to be processed.

One of the reasons for this, naturally, is the ability to do such file processing entirely in RAM. COMMAND.COM must reload a batch file from disk for each line of input, making it necessary to keep the file available throughout its processing. Reaching this file into RAM makes it possible to copy a sequential-based file from a system physically yet continue its execution.

File swapping techniques used in the same area, in which 4DOS has greatly improved on the standard command interpreter. Rather than reloading the same copy that was used to load the program into RAM, 4DOS is able to perform the streamlining use of RAM, including any data storage, if it's not necessary to swap in a backup file. This makes it possible for the program to shrink itself to only 256 bytes in just a DOS RAM area, even on a 386 or higher system that has the ability to load the rest of the program in upper memory.

Version 3.0 of 4DOS and its earlier, N4DOS, used command line switches supplied on the SHELL. The COMMAND.SYS file, however, did not use swapping as done. In version 4.0 of 4DOS, the command line switches are replaced with 4DOS ENVI file.

In this version, though, to take the best of swapping to Extended Memory (EMS), to Expanded Memory (EMM), to disk and to not swapping at all. You can also specify the size of your environment area and of the buffer in which command line history will be saved. Many additional features are controlled by the ENVI file, version 4.0, such as automatically identifying DIR entry types by the color in which the entries are displayed.

Not to be outdone by other command interpreters, a full featured, optional file HELP system that gives you the details. How to use each of the internal commands from the command line. With all the added power, this feature is needed often.

**No Undocumented Features** One of the most startling things about 4DOS is that it is simple, merely 100% compatible with documented variants of DOS and works across all versions of DOS from 2.0 to 5.0. Even 4DOS 4.00 and earlier, since it is a documented feature and took on which COMMAN14.CSV supports it, however, the standard does implement several of them and documents them accordingly. For comparison, in DOS 6.0 reported that a few undocumented parts of DOS be used, but they remain at a minimum.

For instance, 4DOS could stack use of INT 21h as a back door into the command interpreter. In version 3.01, INT 21h was sectioned off to a special command, it is, however, part of the 4DOS code, so that if you did not want to be used to locate the primary shell, just as with the standard command interpreter.

Since Novell's file manager used with NETWARE required INT 21h to operate properly, compatibility is altered. By version 3.05, 4DOS included an extra INR program, SHELL2F, which dupli-

cated the truncations and made the Nowell menu usable. But in version 4.0 SHELL21 went way. It had been fully integrated into 4DOS itself. An INI file setting was added to control how the shell dealt with Nowell's dependencies on undocumented aspects of COMMAND.COM. SHELL21 being derived from version 3.03 still requires SHELL21 to function with the Nowell menu, although this may have changed with the NL7 version.

Here's the (official) description of SHELL21, taken from the on-disk documentation that accompanied the 4DOS version 3.03 release (thanks to Tom Rawson and Rex Cobb, the creators of 4DOS, for letting me use this excerpt):

SHELL21 accomplishes most of the purpose of INT21. It traps calls programs make to intercept 31 and loads a secondary shell to execute them. With one exception, this is not a task programs that use INT21 work properly.

The exception is programs that use INT21 to issue SET commands intended to modify the master environment. Under COMMAND.COM, the INT21 is executed by the primary command processor and therefore explicitly modifies the master environment when a SET command is executed. Since SHELL21 uses a secondary command processor to execute the commands, a SET command will modify the secondary command processor's environment and not the master environment. As of this writing, the only program known to use INT21 to execute SET commands is Personal Revs from Mansfield Software.

SHELL21 starts the command processor to run a secondary shell. The name of the command processor is obtained from the COMSPEC environment variable that was in effect at the time SHELL21 was loaded. SO416.COMSET environment variable setting in effect when SHELL21 is invoked by an INT21 from another program. In general, this is not a serious matter, but remember that if you change COMSPEC during operation of your system, those changes will be noticed by SHELL21 which already the secondary shell. You can set parameters for a 4DOS shell created by SHELL21 using the 4DSHELL environment variable. Use %ENV% for any secondary 4DOS shell.

## Sample Program: An Environment Editor

It's time to put everything together into a nice, simple program that illustrates this chapter's topics and that is also potentially useful. The program, ENVEDIT, is a command editor for the master environment, not the active environment. It does not change the active shell, or the current active environment. ENVEDIT can be used with any command interpreter that respects INT21 for its own code space and that follows COMSPEC conventions, that is, UNIX and OS/2 point to the same segment address. Even if INT21 is not set up, ENVEDIT's core environment capabilities remain available.

In its restriction, ENVEDIT operated only on the master environment. This restriction made it unusable in such common situations as running a DOS session under Windows, so the current version adds a capability to edit the current active environment. You should note, however, that ENVEDIT can't edit the environment within one DOS session under Windows and have your changes take effect in any other session, because each session has its own, unique current active copy, and you need a way to access the copy that Windows uses to create these individual sessions.

ENVEDIT first locates the desired environment block and then locates the specific variable to be edited. If a variable does not already exist, it displays a brief summary of how it is to be used. For listed variables, a variable's current value is shown for the specific environment. You select which environment to edit by means of the option switches -M for master or -A for active. In the absence of any option switch, the master environment is edited if running from DOS. If running under Windows, the active environment is automatically used.

Most of the specific techniques that ENVEDIT uses have already been explained earlier in this paper. The value of this program is that it shows you how to put the pieces together. In addition to ENVEDIT, based on this program, requires the support routines in ENVPKG.ASM and NLS.VIA.ASM plus a new module, FEV.ASM, to support Turbo C's assembly code.

ENVEDIT (Figure 10-16) has been tested only with Borland C++, but it should work equally well with Microsoft's computers. FEV.ASM has been tested with Turbo Assembler and with MASM V3.1. The revised program has been tested with both COMMAND.COM and 4DOS.COM as the main shell, as well with DOS versions ranging from 3.2 to 6.0.

Figure 10-16 ENVEDIT.C

```

.....
* ENVEDIT.C - Editor for Master Environment Variables *
* Jim Kyle, July 8, 1990 *
* major revision August 22, 1992 - jk *
* *
* gcc envedit.c env.obj envpkg.obj nstovar.obj *
* or cl envedit.c +RMS env.asm envpkg.asm nstovar.asm *
* or bcc envedit.c env.asm envpkg.asm nstovar.asm *
* *
.....
#include <stdio.h>
#include <stdlib.h>
#include <conio.h>
#include <string.h>

#include "undocdos.h" /* defines structures */

#define FP_SEG
#define FP_SEG(i) (*((unsigned *)&i) + 1)
#undef FP_SEG

#define FP_OFF
#define FP_OFF(i) (*((unsigned *)&i))
#undef FP_OFF

extern char far * nstovar( char far * wptr );
extern char far * wstenvp( void );
extern void msg_xy( int *x, int *y );
extern int isWin3( void );
extern int col( void );
extern int row( void );
extern void setcl( int r, int c );
extern int envsiz( char far * envptr );

char far * wenvp; /* perm ptr to environment */
char far * wenv; /* pointer to current var */
char far * wnext; /* pointer to next var */
char far * lstbyt; /* adr of last byte used */
char wname[512], *txtptr, /* working buffer and ptr */
int nmlen, insmode = 0, /* name len, insert flag */
max_x, max_y, /* screen limits */
i, c, /* scratchpads */
begrow, begcol, /* cursor, start of text */
currow, curcol, /* current loc */
endrow, /* end of text */
editing, /* loop control flag */
i_cur, i_max, /* cur, max i in txtptr */
free_env; /* bytes free in env */

void findvar( char * varnam ) /* find var, set txtptr */
{ nmlen = strlen( varnam );
txtptr = NULL; /* present not-found flag */

```

```

while ( *menu )
{ rest = nextvar( menu ), /* "rest" always next one */
  sprintf( vname, "%fs", menu ),
  if( vname[nmlen] == '\0' ) /* possible match found */
  { vname[nmlen] = '\0';
    if ( strcmp( vname, varnam ) == 0 )
    { txtptr = &vname[nmlen+1];
      vname[nmlen] = '\0';
      return; /* found it, get out now */
    }
  }
  menu = rest; /* try again with next */
}

void calcscr( void ) /* calc currow, curcol */
{ begrow = endrow - (i_max / max_x );
  if ( (i_max % max_x) == 0 )
  begrow++;
  begcol = 0;
  currow = begrow + (i_cur / max_x );
  curcol = begcol + (i_cur % max_x );
}

void show_var( void ) /* display var content */
{ setrc( begrow, begcol ); /* set to start */
  printf( txtptr ); /* show the string */
  endrow = row(); /* update end row if scroll */
  if ( ! col() ) /* adjust for line scroll */
  endrow += (max_y-1)
  endrow--;
  calcscr(); /* establish cursor posn */
}

void do_del( void )
{ for ( i=i_cur, txtptr[i], i++ ) /* slide over one to left */
  txtptr[i] = txtptr[i+1],
  if ( i_max && i_cur >= --i_max ) /* decr length */
  i_cur = (i_max - 1), /* and adjust if needed */
  free_env++, /* account for freed byte */
  setrc( begrow, begcol ), /* re-display the string */
  printf( txtptr ),
  endrow = row(); /* hold ending point */
  if ( ! col() ) /* adjust for line wrap */
  endrow--, /* " " now in col 0 */
  putchar( ' ' ); /* erase garbage char */
  calcscr(); /* establish cursor posn */
}

void dochr( void )
{ if ( free_env < 3 ) /* just beep if no space */
  { putchar( 7 );
    return,
  }
  if ( insmode ) /* open up a hole for new */
  { if ( --free_env < 3 ) /* decr freespace count */
    { putchar( 7 ); /* and if too little is */
      return, /* left, beep and quit */
    }
    for ( i = ++i_max; i > i_cur; i-- )
    txtptr[i] = txtptr[i-1],
  }
  txtptr[i_cur++] = (char) c, /* put char down */
  if ( i_cur >= i_max ) /* check for extending it */
  { txtptr[ ++i_max ] = '\0'; /* set new EOS */

```

```

    if ( --free_spc < 3 )           /* decr freespace count */
        { putchar( 7 );           /* if too little is left,
          return;                  /* beep and quit */
        }
    }
    show_var(),                    /* re-display the string */
}

int editx( void )                 /* read kbd, do editing */
{ int retval;
  begrow = row(),
  begcol = 0,
  i_max = strlen( tstrptr );     /* set buffer index limit */
  i_cur = 0;                     /* and current index val */
  show_var();                    /* display the string */
  for ( editing=1, editing, )    /* main editing loop here */
  { setrel( 0, 70 );            /* status message loc */
    printf( MODE: "%s", insmode ? "INS" : "REP" );
    setrel( currow, curcol );    /* keep cursor posn curr */
    switch( c = getch() )
    { case 0:                   /* function key or keypad */
      switch( getch() )
      { case 30:                /* Alt-A, redisplay */
        show_var(),            /* re-display the string */
        break,
        case 32:                /* Alt-B, delete variable */
        printf( "\nDELETE this variable (Y/N)? ");
        if( ( getch() & 80 ) == 89 ) /* 89 = 'Y' */
            { vname[0] = '\0';
              retval = 1;
              editing = 0;
            }
        break,
        case 71:                /* home, goto first char */
        i_cur = 0,
        calccrsr();            /* establish cursor posn */
        break,
        case 72:                /* up arrow */
        if ( ( i_cur - max_x ) > 0 )
            i_cur -= max_x,
            calccrsr();        /* establish cursor posn */
        break,
        case 75:                /* left arrow */
        if ( i_cur > 0 )
            i_cur--,
            calccrsr(),        /* establish cursor posn */
            break;
        case 77:                /* right arrow */
        if ( i_cur < i_max )
            i_cur++,
            calccrsr(),        /* establish cursor posn */
            break,
        case 79:                /* end, goto last char */
        i_cur = i_max,
        calccrsr(),            /* establish cursor posn */
        break,
        case 80:                /* down arrow */
        if ( ( i_cur + max_x ) < i_max )
            i_cur += max_x,
            calccrsr(),        /* establish cursor posn */
            break,
        case 82:                /* insert, toggle flag */
        insmode = !insmode,
        break,
        case 83:                /* delete, remove 1 char */

```



```

        dn_del(),
        break;
    }
    break;
case 8:
    if (!_cur)
        {
            _cur = 1;
            do del();
        }
    break;
case 13:
    retval = 1;
    editing = 0;
    break;
case 27:
    retval = 0;
    editing = 0;
    break;
default:
    if (c >= ' ' && c < 127)
        dochar();
    else
        putchar( '?' );
}
}
setrl( endrow, 0 );
return (retval);
}

void putenvbak( void )
{
    char * locptr;
    int save_size;

    save_size = FP_Off( lstbyt ) - FP_Off( rest ) + 1;
    locptr = (char *)malloc( save_size );

    for( i=0; i<save_size; i++)
        locptr[i] = rest[i];
    for( i=0; vname[i], i++)
        *menu++ = vname[i];
    if( vname[0])
        *menu++ = '\0';
    for( i=0; i<save_size; i++)
        *menu++ = locptr[i];
    free( locptr );
    printf( "\nENVIRONMENT UPDATED.\n" );
}

void doedit( char * varnam )
{
    printf( "Editing '%s'\n", varnam );
    menu = menuup;
    free_env = envsize(menu) << 4;
    findvar( varnam );
    for( stbyt=menu; *lstbyt; )
        lstbyt=nextvar(lstbyt);
    if( lstbyt[1] == 1 && lstbyt[2] == 0 )
        {
            lstbyt += 3;
            while (*lstbyt)
                lstbyt++;
        }
    lstbyt++;
    free_env -= FP_Off( lstbyt );
    if ( !txtptr == NULL )
        {
            free_env -= (nrlen-1);
            if ( free_env < 3 )
                {
                    printf( "Not found, no room to add.\n" );
                    return;
                }
        }
}

```

```

    )
    printf( "Not found; add it (Y/N)? " );
    if( ( getch() & 89 ) != 89 ) /* 89 = 'Y' */
        return;
    for ( i=0; i<nlen; i++ ) /* force to uppercase */
        vname[i] = (char) toupper( vname[i] );
    vname[nlen] = '\0'; /* add the equals sign */
    vname[nlen+1] = '\0'; /* make content empty */
    txtptr = &vname[nlen+1]; /* set text pointer to it */
    putchar( '\n' ); /* start on fresh line */
    inmode = 1; /* and in ISB mode */
}
printf( "Free environment space = %d bytes\n", free_env );
if ( edstat() ) /* do the editing now */
    putenvbak(); /* copy to master env */
else
    printf( "\nENVIRONMENT NOT CHANGED\n" );
    putchar( '\n' );
}

void showvars( void ) /* prints usage message */
{ puts( "USAGE: ENVEDIT [env] varname [(name2) ... 3]" );
  puts( "where varname is the name of an env variable", );
  puts( "and name2, etc., are optional added names", );
  puts( "options: -switch [env] is 'M' to use Master copy", );
  puts( "or 'A' to use Active copy of environment.", );
  puts( "current variable names are: ");
  menu = menuv;
  for ( i = 0; i < 8; i++ )
      vname[i] = '\0';
  while ( *menu ) /* get and print names */
      ( sprintf( vname+8, "%s", menu ),
        for ( i = 8; vname[i] != '\0'; i++ )
            /* all done by for() */
            vname[i] = '\0',
            puts( vname ),
            menu = nstever( menu );
  puts( "Re-run with name(s) of variable(s) to be edited." );
}

char far * actenvp( void ) /* trace back to active */
{ WORD parent, self;
  self = _pdp; /* start with current */
  parent = PARENT( self ); /* prog's parent, */
  do { self = parent; /* and go until shell hit */
    } while ( ( parent = PARENT( self ) ) != self );
  return ( char far *)PRG_P( ENV_FR_PSP( parent ), 0 );
}

void main ( int argc, char **argv )
{ int i;
  menuv = isWin3() ? actenvp() - mstenvp(), /* set default */
  for ( i=1; i < argc; i++ ) /* first check for args */
      { switch( *argv[i] )
        {
          case '-': /* process option switches */
          case '/':
            switch( *(argv[i]+1) )
            {
              case 'M': /* use ZE to get master env */
              case 'm':
                menuv = mstenvp(),
                break;
              case 'A': /* back up to active env */
              case 'a':
            }
        }
      }
}

```

```

    menvp = actenvpf);
    break;
  default
    printf( "Unknown option '%s' ignored\n", argv[i] );
  }
  for( j = 0; j < argc; j++) /* remove from arg list */
    argv[j] = argv[j+1],
  argc =
  /* and reduce counts to fit */
  --;
  break;
}
}

if (argc < 2)
  showvars(), /* list all vars to CRT */
else
  { printf( "Changes will modify %s environment at %p\n",
    menvp = actenvpf( "MASTER" "ACTIVE" ), menvp ),
    new_xy( &max_x, &max_y ), /* set up screen limits */
    while ( --argc ) /* process all vars named */
      doedit( ++argv ),
  }
}

```

The program's last action is to establish a default environment based on the return from `ISWin3`. If the `Y` processes a `Y` option switches found on the command line. When all arguments have been checked, `main` determines what to do next.

If the only thing left on the command line is the program name itself, `argc = 2`, the `showvars` procedure is called. If one or more additional arguments were present, the `doedit` function would be called for each of them in return until all have been processed. In either case, `menvp` points to the environment area that will be listed or edited.

The `showvars` procedure scans the environment using `menvp` and cycles through it using `INTEXT`, displaying each word of each record. The `SET` internal command does almost the same thing, but you are not working in the primary shell. `SET` always deals with the local copy of the environment, rather than with the master. `INSERT` lets you choose between the two in this case.

The `doedit` function searches the variable's environment for a variable again using `INTEXT` to start each search at the beginning of a current `env` block, searching a way the variable tree can be to the top-size index of the current `env` block using `envsz`, so that the amount of free space can be calculated later. The function `doedit` has what to do the actual search.

The `findvar` procedure sets up pointer `isptr` either to the address of the variable's contents in global work, the variable's pointer set, or to `NULL`. The `findvar` checks whether the search was successful. Before checking the real `doedit` does not move another pointer, but just moves a certain register back one past the shell's active garbage if one is present, normally `more`. When this is done, the offset portion of `isptr` is set to the total number of environment block bytes in use, subtracting it from the total size of the block leaves in `freeenv` as the remainder, the number of bytes still free for use.

When `freeenv` is present, `doedit` then checks the search result. If the name is not found, the procedure asks you what to add it as a new variable, first verifying that there's room to do so and correcting the `freeenv` value to account for the name itself. If you do want to add the new variable, the name is copied into the global work buffer `varname`, the equal sign, and a terminating `0` are added. `isptr` is set to the address of the `0` character; the presently empty new contents of `isptr` is used, `flag` is set to indicate `INSERT` mode operation. Otherwise, `doedit` returns to `main` to process the next variable of the input list.

If the variable was found, or if a new one is to be added, `doedit` reports the number of bytes still free to use and calls `edit` to do the actual editing. If `edit` returns a `nonzero` value, it means

—completion of the editing operation—`putenvbak` is called to copy the work buffer `vname` back into the environment block, sliding other data around as necessary to make things fit. If the value returned by `editvar` is zero, indicating an F5 key bailout, the message "ENVIRONMENT NOT CHANGED" is displayed and `dosedit` returns.

—more—we look at `editvar`, which contains the bulk of ENVEDIT's complexity, let's see how `findvar` searches for the name. The three pointers that address the block `mem`, `rest` and `stbyte` are adjusted by `var`'s arguments; `var` is global in scope, so that all procedures in ENVEDIT can use them. `findvar` contains `getnextpointer`, actually one of those in the `args` array, as its argument. `findvar` first calls `findvar` to locate and the `mem` pointer is initialized to the value of `mem`, the permanent pointer to the block to be edited.

Each time that `findvar` is entered, the pointer `stbyte` is set to NULL and the global variable `mem` is set to the long list `findvar`'s argument. Then `findvar` goes into a loop that continues until either the end of the variable set is reached, when the byte pointed to by `mem` is zero, or the argument is found as a variable name.

Without copying the pointer, `rest` is first set to the address of the next variable, using `nextvar`. The cursor compares `var` then `rest` from the environment block to the global work buffer `vname`, by means of `strcmp`, `strlen`, `strcpy`, `isalnum`, and its `is` format modifier. Next, the character at position `rest` in the work buffer is checked; if it is an equal sign, the variable's name is the right length to be possibly a match; if not, no match is possible, so no time is wasted trying to compare the strings.

If the strings split, the equal sign is temporarily replaced by 0, and the `strlen` function string is used to compare the `var` and the argument without regard to case. If they match, the equal sign is put back, `stbyte` is set to the last byte after it, and `findvar` returns successfully. If the strings differ, or the strings fail to match, `rest` is copied to `mem`, and the process repeats for the next variable in the master block.

If we successfully return from `findvar`, `mem` still points to the first byte of the variable being edited; the work buffer contains an exact copy of the entire variable, including its name and the trailing null sign—separator—and `stbyte` points to the first byte of its value in `vname`. These facts are critical to the operation of `editvar`, which does the actual editing.

Of course, if `findvar` variables are set to parse the cursor position and `showvar` is called to display the variable on the screen, the main loop's then centered; control remains within this loop until you press Alt+F, Enter, or Esc. Any of these keys causes control variable editing. Enter and Alt+D set the cursor value to 1, and Esc sets the `return` value to 0. The cursor is then positioned just past the end of the variable on the screen, and control returns to `dosedit`.

If `return` is 0, `showvar` is called; it positions the cursor to the saved starting position, displays the entire contents of the variable using `fprintf`, saves the ending cursor position, and calls the function `calccur` to calculate the necessary adjustments to the saved starting position and to establish the current cursor position within the variable.

If Alt+F is struck, simply displays the display window; anything cause it to become confused. Alt+D causes some ENVEDIT set you want to delete the variable entirely, rather than just change it. What you have selected, you really mean it; the program zeroes the first byte in `vname` and `return` is zero; this means those that occur when you press Enter. This way, no bytes move back to the master block, but each of the variable's address one up the day.

If `return` is 1, `calccur` is called, as would be expected, with the restriction that the cursor cannot move out of the variable. `calccur`'s operation has no effect if the cursor is on the top line of the variable, left arrow keys moving it to the cursor's so the first character, and so forth. The Home and End keys move the cursor to the first and last characters, respectively, of the variable. The code for all six of these keystrokes simply changes the value of `cur` as appropriate, then calls `calccur` to do the heavy work.

Once `edit` has completed and returns a nonzero value, `putenvbak` handles the job of moving the edited variable back into the master block. It does this by first calculating the number of bytes that must be saved using the values of `rest` and `before` that were set up earlier, and then copying that material from the master block to a temporary block of memory obtained using `malloc`.

The final content of `envp[4]` is then copied back to the specified environment block, starting at the address indicated by `envp`. Finally, the saved material is copied in following the unaltered data and the temporary block released by `free`. The final action of `putenvbak` is to display or save the line "ENVIRONMENT UPDATED."

The module `ENVASM` contains `ISWMB` and functions that are used when you are interactively editing the environment variables. It is not shown here, but you can find it on the company disk etc. To some extent, the functions in `ENVASM` duplicate others found in the Borland and Microsoft libraries, but the library functions differ greatly between the compilers, so whether or not these video routines simplify the rest of the program by providing identical operation regardless of the compiler you choose.

Here's what the operation of our revised `ENVEDIT` looks like. First, we just tell it to edit variable `OLDPATH`:

```
D:\DOS2\CHAP10> envedit oldpath
Changes will modify MASTER environment at D514:0000
Editing 'oldpath'
Free environment space = 270 bytes
D:\WIN31, D:\VB, D:\UTILS, D:\DOS, D:\TAPE, C:\UV, C:\VMAX6, C:\ADDSOR
ENVIRONMENT NOT CHANGED
```

The second line indicates that `ENVEDIT` has defaulted to editing the master environment if the first line indicates that I pressed `F`, rather than `A`, to leave the program and so the environment didn't change. Next, I repeated the command, but added the `/a` option switch to specify the active environment:

```
D:\DOS2\CHAP10> envedit /a oldpath
Changes will modify ACTIVE environment at D4C8:0000
Editing 'oldpath'
Free environment space = 262 bytes
D:\WIN31, D:\VB, D:\UTILS, D:\DOS, D:\TAPE, C:\UV, C:\VMAX6, C:\ADDSOR
```

With the content of `OLDPATH` displayed, I used the `DEL` key repeatedly to erase it, to arrive at `D:\WIN31`, then I pressed Enter to change the variable in memory:

```
D:\VB, D:\UTILS, D:\DOS, D:\TAPE, C:\UV, C:\VMAX6, C:\ADDSOR
ENVIRONMENT UPDATED.
```

Repeating the original command sequence shows that no change occurred in the master environment, that's just what should have happened:

```
D:\DOS2\CHAP10> envedit oldpath
Changes will modify MASTER environment at D514:0000
Editing 'oldpath'
Free environment space = 270 bytes
D:\WIN31, D:\VB, D:\UTILS, D:\DOS, D:\TAPE, C:\UV, C:\VMAX6, C:\ADDSOR
ENVIRONMENT NOT CHANGED
```

Finally, adding the `/a` switch at the end of the last list to show that my position makes no difference, gets a report that, unlike the earlier change, was saved to the active environment:

```
D:\DOS2\CHAP10> envedit oldpath /a
Changes will modify ACTIVE environment at D4C8:0000
Editing 'oldpath':
Free environment space = 271 bytes.
D:\VB, D:\UTILS, D:\DOS, D:\TAPE, C:\UV, C:\VMAX6, C:\ADDSOR
ENVIRONMENT NOT CHANGED
```





- CX** size of returned data  
**EA** number of times DPMI interpreter invoked since parameters were last set

**Notes:** The DPMI interpreter was apparently never released due to its large memory requirements and slow operation.

If the given buffer size is at least two bytes but less than the size actually used by the parameter table, the first CX bytes will be copied into the buffer but no error will be returned.

#### Format of parameter table:

Offset	Size	Description
00h	WORD	size in bytes of following data
02h	WORD	max # of number of error corrections per line to attempt
04h	WORD	order in which to correct errors: 00h left to right 01h right to left 02h most serious first

#### —DOS 2.70—

00h	BYTE	unknown
07	DWORD	pointer to DPMI statistics record (see INT 2B AH 03h)

#### —DOS 2.71—

00h	WORD	unknown
08h	DWORD	pointer to DPMI statistics record (see INT 2B AH 03h)

Note: A statistics record may be placed in a user buffer by changing the pointer with INT 2B AH 02h.

See Also: INT 2E AH 04h INT 2Bh AH 02h INT 2Bh AX 000h

The parameter table is the standard data exchange mechanism of DOS for which the function is valid on this set of systems (2.70 and 2.71 only). Various other variations are also used: DOS 2+ indicates that the carry flag is set; all known versions of DOS treat 2.00 and 00 as binary, and DOS 3.1-3.3 indicates that the interrupt is valid for DOS versions 3.1 through 3.3 inclusive.

Various interrupt functions are called "internal" by DOS and should be implemented by a user program rather than a utility. To further distinguish these functions, "Called with" is used instead of "Called with".

"Called to" means that the register or register pair contains the address of the indicated item rather than the item itself.

Because descriptions have an order of data, that the register only applies for the value indicated. If a procedure is defined by register description. For this example, the value in AX is meaningful only if the carry flag is set; otherwise, the value of AX is completely unimportant. The carry flag is clear.

Italicized text indicates that the information is not entirely certain and that particular care—even attention to details—should be exercised when attempting to use it.

A Note at the end of some descriptions refers to the function in general or to the description of a data structure. Note that only one field (not a byte) Return section apply to the function in general, while those that apply to the data structure description apply only to the data structure (see below).

Many data structures are defined by DOS, as between versions of DOS. To save space, the different variations are combined into one description; the fields are assumed to be common to all variations unless otherwise specified. For example, the first three fields are common to both 2.70 and 2.71. In DOS 2.70, the fourth field is a byte, while in DOS 2.71 it is a word, shifting the remaining field.

One or more fields in a data structure may be pointers to additional data structures. Such data structures are often described in other entries of the appendix; in this case, under Interrupt 2Bh function 03h.







10h	DWORD	file size
14h	WORD	date of last write (see 2f 5700h)
16h	WORD	time of last write (see 2f 5700h + 10x41h)
18h	8 BYTES	reserved (see below)
20h	BYTE	record within current block
21h	DWORD	random access record number. If record size is 64 bytes, high byte is omitted

**Note:** To use an extended LFB you must specify the address of the LFB flag, if offset -7 rather than the address of the drive number field.

#### Format of reserved field for DOS 1.0:

Offset	Size	Description
16h	WORD	location of directory. If high byte 11h, low byte is logical ID
18h	WORD	number of first cluster in file
1Ah	WORD	last cluster number accessed (a volume)
1Ch	WORD	entry of relative cluster number within file. 0 = first cluster of file
1Eh	BYTE	dirty flag (00h=not dirty)
1Fh	BYTE	unused

#### Format of reserved field for DOS 1.10 1.25:

Offset	Size	Description
18h	BYTE	bit 7: set if logical device bit 6: not dirty bits 5-0: disk number or logical device ID
19h	WORD	starting cluster number on disk
1Bh	WORD	current absolute cluster number on disk
1Dh	WORD	current relative cluster number within disk
1Fh	BYTE	unused

#### Format of reserved field for DOS 2.x:

Offset	Size	Description
18h	BYTE	bit 7: set if logical device bit 6: set if open bits 5-0: unknown
19h	WORD	starting cluster number on disk
1Bh	WORD	unknown
1Dh	BYTE	unknown
1Eh	BYTE	unknown
1Fh	BYTE	unknown

#### Format of reserved field for DOS 3.x:

Offset	Size	Description
18h	BYTE	checksum of system file table entry for file
19h	BYTE	attributes
		bits 7,6: 00 SHARE.EXE not loaded, disk file 01 SHARE.EXE not loaded, character device 10 SHARE.EXE loaded, remote file 11 SHARE.EXE loaded, local file
		bits 5-0: low six bits of device attribute word

#### SHARE.EXE loaded, local file (DOS 3.x and 5.0):

1Ah	WORD	starting cluster of file on disk
1Ch	WORD	DOS 3.x: offset within SHARE if sharing record (see 2f 526)

DOCS 5.0: unique sequence number of sharing record

1Fh BYTE file attribute

1Hh BYTE *unknown*

**SHARE EXE loaded, remote file—**

1Ahh WORD number of sector containing directory entry

1Ch WORD relative cluster within file of last cluster accessed

1Eh BYTE absolute cluster number of last cluster accessed

1Hh BYTE *unknown*

**SHARE EXE not loaded—**

1A STARTING CLUSTER OF FILE (DOCS 5.0) OR SECTOR (DOCS 4.0)

1Bh WORD starting cluster of file

1Eh WORD number of sector containing directory entry

1Hh BYTE number of directory entry within sector

**Note:** For DOS 4.0, the starting cluster of file is stored at the address of the device driver header, and then the BYTE at 1Ah is overwritten.

**INT 21h Functions 18h, 1Dh, 1Eh**

**DOS 1+**

**NULL FUNCTIONS FOR CP/M COMPATIBILITY**

In CP/M, the functions INT 21h 18h, 1Dh, and 1Eh were null functions. In MS-DOS, an equivalent to function 1Eh was reintroduced with DOS 2.0.

Call With

AH 18h, 1Dh, or 1Eh

Returns

AX 00h

**Note:** Function 18h corresponds to the CP/M BIOS "get bit map of floppy drives" function (DOS 4.0 equivalent is INT 21h 1Eh). Function 1Dh corresponds to the CP/M BIOS "set bit map of floppy drives" function (DOS 4.0 equivalent is INT 21h 1Eh).

See Also: 21h 20h 21h 4459E

**INT 21h Function 1Fh**

**DOS 1+**

**GET DRIVE PARAMETER BLOCK FOR DEFAULT DRIVE**

Returns the address of a disk description table for the current

Call With

AH 1Fh

Returns

AX status  
 00h: success

DS:BX: pointer to Drive Parameter Block (DPB) (see below for DOS 1.x, 2.1, 2.0h for DOS 2.x)

1Eh: invalid drive

**Note:** This call was undocumented prior to the release of DOS 5.0. In newer, only the DOS 4+ version (DPB) is documented.

See Also: 21h 32h

**Format of DOS 1.1 and MS-DOS 1.2S drive parameter block:**

Offset	Size	Description
00h	BYTE	sequential device ID
01h	BYTE	logical drive number (0-5)
02h	WORD	bytes per sector

04h	BYTE	highest sector of directory entry
05h	BYTE	shift count to convert cluster to sectors
06h	WORD	starting sector of first FAT
08h	BYTE	number of copies of FAT
09h	WORD	number of files in current
0Ah	WORD	number of first data sector
0Bh	WORD	high water mark of sectors of fat clusters + 1
0Ch	BYTE	sectors per FAT
10h	WORD	starting sector of directory
12h	WORD	address of allocation table

**Note:** The DOS 1.0 table is the same except that the first 7 of last fields are missing (see 21-32 for the DOS 2+ version).

## INT 21h Function 20h

DOS 1+

### NULL FUNCTION FOR CP/M COMPATIBILITY

This function corresponds to the CP/M BIOS function "get/set default user (sublibrary number)", which is normally used under MS-DOS.

Call With

AH = 20h

Returns

AH = 4Fh

See Also: 21-10h, 21-4459h

## INT 21h Function 26h

DOS 1+

### CREATE NEW PROGRAM SEGMENT PREFIX

Although undocumented, the PSP contains undocumented fields.

Call With

AH = 26h

DX = segment of memory to create new PSP. See below for format.

**Note:** The new PSP is updated with memory size information. INTs 22h, 23h, and 24h are taken from the current interrupt vector table.

DOS 2+ assumes that user has a valid segment for the PSP group.

See Also: 21-44, 21-50, 21-5D, 21-5E, 21-62, 21-67h

### Format of PSP:

Offset	Size	Description
00h	7 BYTES	INT 21h vector for CP/M. All 0 for programs running under DOS. The address of the segment for a new PSP.
02h	WORD	segment of parent program. (Reserved to program.)
04h	BYTE	unused filler
05h	35H	CP/M Access Mode (always 0) 00000000
06h	WORD	BLKs (DOS 2+ PSPs are 16 to 2 <sup>14</sup> 4h) (see 4000h)
08h	2 to 11h	CP/M segment of program (segment of COM files reserved for MIP/MS-DOS)
0Ah	3 WORDS	segment of MIP file (segment of COM)
0Ch	WORD	segment of MIP file (high address)
12h	3 WORDS	DOS 2+ segment of parent program (segment of parent PSP)
16h	WORD	DOS 2+ segment of parent program (segment of parent PSP)
18h	70 to 115h	DOS 2+ segment of parent program (segment of parent PSP)
2Ch	WORD	DOS 2+ segment of parent program (segment of parent PSP)



**INT 21h Function 30h****DOS 2+****GET DOS VERSION**

Although documented, the OEM number this function returns are not documented.

**Call With**

AH 30h

**—DOS 5.0—**

AI what to return in BH  
 00h OEM number (as for DOS 2.0-4.0)  
 01h version flag

**Returns:**

AL major version number (00h if DOS 1.x)  
 AH minor version number  
 BH CX 24-bit user-selectable OEM number (most versions do not use this)

**—if DOS < 5 or AL=00h—**

BH OEM number  
 00h IBM  
 05h Zenith  
 16h DEC  
 23h Olivetti  
 29h Toshiba  
 40h Hewlett-Packard  
 99h STARC (H) a collection of OEMs (DOS NETWORK DOS SMP DOS)  
 EAh Microsoft Phoenix

**—if DOS 5.0 and AL=01h—**

BH version flag  
 bit 3: DOS is in ROM  
 other: reserved 0

**Notes:** The OS 2.1.x (original) BIOS returns major version 0Ah (10) the OS 2.2.x (coming at last) returns major version 14 (20) and later versions of the Windows NT BIOS have returned major version 15 (30) (NT BIOS also includes a DOS 5.1N1 24-bit AX 3200h return as DOS 5.50). See Chapter 4.

DOS 4.01 and 4.02 identify themselves as version 4.00.

Compaq MS-DOS 3.31 (or any MS-DOS 3.31) and others identify themselves as a C-DBS by returning OEM number 00h.

The version returned for DOS 4.0x may be modified to entries in the special program list (see 21/52h).

The user-selectable OEM for DOS 5.0x is being modified by SEEVER (see 21/350h) to get the true version number.

See Also: 21/310h; 21/1221h

**INT 21h Function 32h****DOS 2+****GET DOS DRIVE PARAMETER BLOCK FOR SPECIFIC DRIVE**

Determines the address of a disk description table for the specified drive.

**Call With**

AH 32h

DI drive number (00h= default, 01h=A: etc.)

**Returns:**

AI status  
 00h successful

DPB points to Drive Parameter Block (DPB) for specified drive. If it is invalid or network drive.

**Notes:** 1. For OS 2 compatibility, this supports the DOS 3.3 version of the call with the exception of 1 WORD of extra data. This call replaces the DPB's reading the disk; the DPB may be accessed via `DPB->Sector`, `DPB->DiskAccess` or `DPB->Disk`. See Chapter 8.

2. This is a completely undocumented for DOS 5.0, but was well-documented in prior versions. On the DOS 4+ version of the DPB was documented, however.

See Also: 21/1fh, 21/5bh, 21/53h

#### Format of DOS Drive Parameter Block:

Offset	Size	Description
00h	BYTE	drive number (00h-A, 01h-B, etc.)
01h	BYTE	unit number within device (driver)
02h	WORD	bytes per sector
04h	BYTE	highest sector number within a cluster
05h	BYTE	shift count to convert clusters into sectors
06h	WORD	number of reserved sectors at beginning of drive
08h	BYTE	0x16 of FAT
09h	WORD	number of root directory entries
0Bh	WORD	number of first sector containing user data
0Dh	WORD	$(\text{bytes per sector} / \text{bytes per cluster}) \times \text{cluster} + 1$
0Eh	BYTE	16-bit FAT if sector size = 0100h, else 12-bit FAT
0Fh	WORD	sector number of first directory sector
10h	DWORD	address of device driver header
16h	BYTE	media ID byte
17h	BYTE	00h if disk accessed, 11h if not
18h	DWORD	pointer to next DPB
<b>DOS 2.x—</b>		
1C1	WORD	cluster number of start of current directory (0000h if not known, 1111h if not known)
111	64 BYTEs	AM/IF/permissions of current directory for drive
<b>DOS 3.x—</b>		
17	WORD	cluster number of current search for free space when writing
1F	WORD	number of bytes reserved for drive (1111h if not known)
<b>DOS 4.0 &amp; 6.0—</b>		
0Fh	WORD	number of sectors per FAT
11h	WORD	sector number of first directory sector
13h	DWORD	address of device driver header
17h	BYTE	media ID byte
18h	BYTE	00h if disk accessed, 11h if not
19h	DWORD	pointer to next DPB
1Dh	WORD	cluster number of current search for free space when writing (usually the last cluster allocated)
1Eh	WORD	number of reserved sectors on drive (1111h if not known)

#### INT 21h Function 3302h

DOS 3+

#### GET AND SET EXTENDED CONTROL-BREAK CHECKING STATE

Set new state for the extended Control-Break checking flag, and returns its old state.

Call With

AX = 3302h



DI new state (00h for OFF or 01h for ON)

**Returns**

DI old state of extended BREAK checking

**Note:** This function does not use any of the DOS internal stacks and may thus be called at any time even during another INT 21h call. One possible use is modifying Control Break check from within an interrupt handler or pop-up ISR. See Figure 6-5.

**INT 21h Function 34h****DOS 2+****GET ADDRESS OF CRITICAL SECTION (InDOS) FLAG**

Return the address of flag which indicates when code within DOS is being executed. Useful for insuring safe to call DOS functions.

**Call With**

AH 34h

**Returns**

ES:BX pointer to one-byte InDOS flag

**Notes:** In code of DOS is incremented whenever an INT 21h function is called and decremented whenever one completes. However, InDOS alone cannot sufficient to determine when it is safe to enter DOS, as the critical error handling decrements a DOS flag increments the critical error flag for the duration of the error handling. InDOS will be zero if DOS is not in error handling. InDOS will be non-zero if DOS is in error handling. InDOS will be non-zero if DOS is in error handling.

The error flag is located at address 0000:0000. InDOS will be non-zero if DOS is in error handling. InDOS will be non-zero if DOS is in error handling.

The error flag is located at address 0000:0000. InDOS will be non-zero if DOS is in error handling. InDOS will be non-zero if DOS is in error handling.

For DOS 3.1+ the undocumented GET SFA function (21/5100h) can be used to get the address of the critical error flag.

This function is always called prior to the release of DOS 3.0.

The response code of this function is always less than 64h.

See Also: 21/5100h, 21/5100h, INT 20h

**INT 21h Function 3700h****DOS 2+****"SWITCHAR" GET SWITCH CHARACTER**

Display the character which was used to introduce command and file processing as done by the DOS program (version 4.1) and higher. Not recommended because of the state of progress.

**Call With**

AX 3700h

**Returns**

AL status  
00h successful  
DL current switch character  
Fbh unsupported substitution

**Notes:** This function is undocumented in some OEM versions of some releases of DOS and is supported by the OS/2 compatibility box.

The returned value is always 2Fh for DOS 5+.

See Also: 21/3701h

**INT 21h Function 3701h****DOS 2+****"SWITCHAR" SET SWITCH CHARACTER**

Set a new character to be character which is used to introduce command and file processing.

**Call With**

AX 3701h  
 DI new switch character

**Returns**

AI status  
 00h successful  
 FFh unsupported subfunction

**Notes:** This function is broken in some OEM versions of some releases of DOS and is supported by the OS-2 compatibility box.

This call is ignored by DOS 5.

See Also: 21-3700h

**INT 21h Functions 3702h and 3703h****DOS 2.x and 3.3+ only****"AVAILDEV" SPECIFY \DEV\ PREFIX USE**

Get or set the status of a flag which takes a DEV prefix to character device names manually.

**Call With**

AH 37h  
 AI subfunction  
 02h get availdev flag  
 Returns:

DI 00h if DEV must precede storage or device names  
 + success if DEV is optional

03h set availdev flag

DI 00h DEV is mandatory  
 nonzero DEV is optional

**Returns**

AI status  
 00h successful  
 FFh unsupported subfunction

**Notes:** All versions of DOS from 2.00 allow DEV to be prepended to device names without generating an error if the device DEV is not a character device; other paths generate an error if they do not.

Although DOS 3.3+ accepts these calls, they have no effect and subfunction 02h always returns DI 00h.

**INT 21h Function 3Fh****Workgroup Connection****WORKGRP SYS GET ENTRY POINT**

WORKGRP SYS GET ENTRY POINT For Microsoft's Workgroup Connection, which permits communication with a server. Workgroup or Worktop, or FAX Manager networks. This call is one way to determine the address of the API handler.

**Call With**

AH 3Fh  
 BX file handle for device "\NETSHARE PS"  
 CX 0000h  
 DS:DI pointer to buffer to store pointer to C:\MS-A\402h\WORKGRP\SYS"

**Returns**

CF clear if successful  
 AX number of bytes actually read (0 if at EOF before call)  
 CF set on error  
 AX error code: 05h,06h (see AH-59h)

**INT 21h Function 3fh****Driver Names****READ DEVICE**

A number of device names are used by various drivers to provide interfaces to the devices by performing reads or writes rather than IOCTL calls. See the list below for a sampling of special device names and the full interrupt list on disk for details on several of these devices.

See Also: 2f-4402h, 2f-9400h

**Special Device Names:**

03CMPAQ	
386MANS	Quantix 386 to the MAX
CACFCMPQ	
EMMXXXX	Expanded Memory Manager (not all memory managers provide data on access to this device name)
EMMQXXX	QEMM 386 with disabled EMM
EMMXXXX	disabled EMM
NIUSEHPS	Worker 386 (WORKER386) WORKGROUPS
QEMM386S	Quickcheck's QEMM 386
SMARTAAR	SmartDrive 3.5
MMXXXXX	HMM386SYS (not used for an API by HMM386 but possibly by other extended memory managers)
SIdebugDD	Windows low level debugger
SMMXXXXX	disabled EMM
MMXXXXX	disabled EMM

**INT 21h Function 40h****DOS 2+****ADJUST FILE SIZE**

Set the size of the specified file to the current file position as set by SEEK\_21\_42h (truncating or expanding file as necessary). This is a little-known, but documented, feature of the write-to-file call accessed by specifying a value of 0x40h. This feature is a debug cover.

**Call With**

AH	40h
CX	0000h
BX	file handle

**Returns**

CF	clear if successful
AX	0000h
CF	set on error
AX	error code (see 2f-50h)

**BUG** A series of zero bytes to adjust the file size will appear successful when it is the intended file write is successful. If there is not enough disk space for the expanded file, one should therefore check a series of bytes (not extended by seeking to 0 bytes from the end of the file) (21/4202h/0x0, 0x0).

See Also: 2f-1109h

**INT 21h Function 41h****DOS 2+****"UNLINK" - DELETE FILE**

When executed as a 21-5100h, this function has the undocumented behavior of allowing wildcards in the filename and enabling a attribute mask, deleting all files matching the wildcards and attribute mask.

**Call With**

AH 41h  
 DS:DX pointer to AM I/O buffer, not necessarily attribute mask for deletion; server call may see XIO

**Returns**

CF clear if successful  
 AX destroyed; IOPL 43  
 AL appears to be drive of deleted file  
 CX set on error  
 AX error code (02h-03h-05h); see 21-96h

**Notes** 1. In DOS 3.1, wildcards are allowed in the filename; in versions 2.1-5.1000h, in which case the filename must be enclosed in quotes by 21-60h, and only files matching the attribute mask in CX are deleted.

DOS does not erase the file's data; it merely becomes inaccessible because the FAT chain for the file is cleared.

Deleting a file which is currently open may lead to memory corruption. Unless SHARE is loaded, DOS does not close the handles belonging to the deleted file, thus allowing writes to a nonexistent file.

See Also: 21-13h, 21-51000h, 21-60h, 21-1113h

**INT 21h Function 4302h****Chicago****GET VOLUME INFORMATION**

CALL INT 21h, DOS, WinSys 1, 21-402, 21h by the INT 21h equivalent of the Win32 GetVolumeInformation API. The function returns information on the entry to store long filenames, is supported on FAT, whether the entry preserves lower case filenames, the maximum length of a path name, necessary to allow for others; see 21-71, a discussion; see 21-71, 21-72 and chapter 8.

**INT 21h Function 4302h****DR-DOS 3.41+****GET ACCESS RIGHTS**

The function reports on the access rights that can be performed on a specified file without being required to provide a password.

**Call With**

AX 4302h  
 DS:DX pointer to AM I/O buffer

**Returns**

CF clear if successful  
 CX access rights

bit 0	owner delete requires password
bit 1	owner execute requires password (HexOS)
bit 2	owner write requires password
bit 3	owner read requires password
bit 4	group delete requires password
bit 5	group execution requires password (HexOS)
bit 6	group write requires password
bit 7	group read requires password
bit 8	world delete requires password
bit 9	world execution requires password (HexOS)
bit 10	world write requires password
bit 11	world read requires password

AX	equals CX (DR-DOS 5.0)
CF	set on error
AX	error code

**Notes:** This protection scheme has been coordinated in a letter to Digital Research (Novell operating systems DR-DOS 3.41+, DRMDOS 5.x and EtcOS 2+).

Only EtcOS actually uses the "execute" bits (DR-DOS 3.41+ treats them as "read" bits).

This function is documented in DR-DOS 6.0 and corresponds to the "Get/Set File Attributes" function, subfunction 2, documented in Concurrent DOS.

DR-DOS 3.4x, 5.x only use bits 0-3. Only DR-DOS 6.0 using a DRMDOS 5.x security system allowing for users and groups uses bits 4-11.

See Also: 21/4303b

## INT 21h Function 4303h

**DR-DOS 3.41+**

### SET ACCESS RIGHTS AND PASSWORD

Specify which operations may be performed on a file without a password, and optionally set the file's password.

**Call With**

AX	4303h
CX	access rights bits 0-11 see 21/4302h bit 15 new password is to be set
DS:DX	pointer to ASCII path name
DI:Y	new password of CX bit 15 is set. Blank padded to 8 characters

**Returns**

CF	clear if successful
CF	set on error
AX	error code

**Notes:** If the file is already protected, the old password must be added after the path name, separated by a semicolon.

This function is documented in DR-DOS 6.0 and corresponds to the "Get/Set File Attributes" function, subfunction 3, documented in Concurrent DOS.

See Also: 21/4302h, 21/4454b

## INT 21h Function 4304h

**DR-DOS 5.0+**

### GET ENCRYPTED PASSWORD

Get the encrypted password for a file.

**Call With**

AX	4304h
	<i>additional arguments (if any) unknown</i>

**Returns**

CF	clear if successful
AX	unknown
CX	unknown (same as AX)
CF	set on error
AX	error code (see 21/50b)

See Also: 21/4303b, 21/4305h

## INT 21h Function 4305h SET EXTENDED ATTRIBUTES

**DR-DOS 5.0+**

This function permits the extended attributes, and optionally the encrypted password, for a file to be set.

### Call With

AX 4305h  
*additional arguments if any unknown*

### Returns

CF clear if successful  
CF set on error  
AX error code (see 21 50h)

See Also: 21 4304h

## INT 21h Function 4400h IOCTL GET DEVICE INFORMATION

**DOS 2+**

Allows a user to get device information word (a number of undocumented attributes).

### Call With

AX 4400h  
DX handle

### Returns

CF clear if successful  
DX device information word  
character device  
11 3 seconds to process I/O requests (see 21 4405h-4406h)  
13 output until busy supported  
11 0 = OPEN\_CLOSE calls (see 21 4406h)  
set indicates device

- 0 EOF
  - 1 0 = 0
  - 4 device is special (uses INT 29h)
  - 4 clock device
  - 2 NFI device
  - 1 standard out
  - 0 standard input
- attributes
- 1 0 = 0 (DOS 3)
  - 0 = 0 (DOS 5)
  - 0 = 0
  - 5 0 = 0 (INT 24h) disk space on write
  - \* clear indicates file
  - 0 file has not been written
  - 5 0 = drive number (0-A)

AX destroyed

CF set on error  
AX error code (01h/05h/06h) (see 21 50h)

**Note:** The value returned in the device information word if the handle refers to a character device.

See Also: 21 4401h, 21 52h, 2F 122Bh









10h single tasking

14h multitasking

**AI** operating system version ID (see below for single tasking, 21-4451h for multitasking)

**DX** same as AX

**Note** For DR-DOS 5.00, the stored version ID is 00h for single tasking.

Use the version ID to determine the multitasking capabilities of 4451h for multitasking, or 4452h for single tasking. See the "Tasking Capabilities" section of DR-DOS 5.0 Chapter 4.

See Also: AX-4412h AX-4451h AX-4459h

#### Values for version ID

00h DR-DOS 5.00

01h DR-DOS 3.41

02h DR-DOS 3.42

03h DR-DOS 5.00

04h DR-DOS 6.00

05h EGA-DOS

06h DR-DOS 6.00 March 1993 update for Win

7Fh Novell-DOS 4.0

### INT 21h Function 4454h

**DR-DOS 3.41+**

#### SET GLOBAL PASSWORD

Specify the master password for accessing files.

**Call With**

**AX** 4454h

**ES:BX** pointer to password (max. length 16 characters)

See Also: 21-4303h 21-4414h

### INT 21h Function 4456h

**DR-DOS 5.0+**

#### HISTORY BUFFER CONTROL

**Call With**

**AX** 4456h

**DI** flag

bit 0 set for COMMAND.COM history buffer; bit 1 for application

**Returns**

**AX**

20h if DI bit 0 set; 30h if clear; DR-DOS 6.0

**Note** For DR-DOS 5.0, the COMMAND.COM or DR-DOS 6.0

### INT 21h Function 4457h

**DR-DOS 5.0+**

#### SHARE/HILOAD CONTROL

Share/hi-load control. The SHARE/SHR or HILOAD/HLI command should place the DR-DOS kernel into high memory.

**Call With**

**AX** 4457h

**DH** subfunction

00h enable/disable SHARE

01h new state

00h disable

01h enable

else Returns AX unknown

*01h* get HHCAD status

Returns AX status  
0000h on  
0001h on

*02h* set HHCAD status

DI new state (00h off, 01h on)

Returns AX unknown

other

Returns AX unknown

**Note** This function is called by COMMAND.COM of DR-DOS 6.0 if an exported file is supported on a file system of DR-DOS. Note: DR-DOS

See Also: AX-4457h DX-FFFFh

## INT 21h Function 4457h Subfunc FFFFh

DR-DOS 6.0

### GET SHARE STATUS

Call With

AX 4457h

DX FFFFh

Returns

AX SHARE status

See Also: 21-1000h

## INT 21h Function 4458h

DR-DOS 5.0+

### GET POINTER TO TABLE OF VARIOUS INTERNAL VALUES

Determine the address of an internal table of useful DR-DOS values and pointers.

Call With

AX 4458h

Returns

ES:BX pointer to internal table (see below)

AX *value in 01Sub of DR-DOS 5.0 03Sub of DR-DOS 6.0*

See Also: 21-4457h

#### Format of internal table:

Offset	Size	Description
00h	DWORD	pointer to unknown item
04h	DWORD	unknown
08h	WORD	K of extended memory at startup
0Ch	BYTE	number of far jump entry points
0Eh	WORD	segment containing the points to DR-DOS entry points (see below)
<b>—DR-DOS 6.0—</b>		
10h	WORD	0000h if valid; 0 if HMA (00000000H) is free
		HMA memory block (see below) or 0000h if none; 00000000H if
12h	WORD	pointer to segment table if segment addresses set in
		CONFIG or 0000h if already used
14h	WORD	0 if kernel loaded in HMA (offset in HMA or first use)
		HMA memory block (see below) or 0000h if none; 00000000H if

**Note** The segment used for the DR-DOS 6.0 CONFIG table in these variables, excluding COMSPEC, VER, and OS, is not useful for programs, drivers called from CONFIG.SYS, the valid set to zero later and the area lost.

**Format of jump table for DR-DOS 5.0-6.0:**

Offset	Size	Description
00h	5 BYTES	far jump to kernel entry point for CP/M CALL 5
05h	5 BYTES	far jump to kernel entry point for INT 20
0Ah	5 BYTES	far jump to kernel entry point for INT 21
0Fh	5 BYTES	far jump to kernel entry point for INT 23 <i>REIF</i>
14h	5 BYTES	far jump to kernel entry point for INT 24
19h	5 BYTES	far jump to kernel entry point for INT 25
1Eh	5 BYTES	far jump to kernel entry point for INT 26
23h	5 BYTES	far jump to kernel entry point for INT 27
28h	5 BYTES	far jump to kernel entry point for INT 28
2Dh	5 BYTES	far jump to kernel entry point for INT 2A (IRET)
32h	5 BYTES	far jump to kernel entry point for INT 2B (IRET)
37h	5 BYTES	far jump to kernel entry point for INT 2C (IRET)
3Ch	5 BYTES	far jump to kernel entry point for INT 2E (IRET)
41h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
46h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
4Bh	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
50h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
57h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
5Ch	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
61h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
68h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
6Dh	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
72h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
79h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
7Eh	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
83h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
8Ah	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
8Fh	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
94h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
9Bh	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
A0h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
A7h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
AC	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
B1h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
B8h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
BFh	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
C6h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
CDh	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
D4h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
DBh	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
E2h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
E9h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
F0h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
F7h	5 BYTES	far jump to kernel entry point for INT 2F (IRET)
FEh	5 BYTES	far jump to kernel entry point for INT 2F (IRET)

**Note:** All of these entry points are redirected through this jump table to allow the kernel to be relocated into high memory while leaving the actual entry addresses in low memory for maximum compatibility.

**Format of HMA Memory Block (DR-DOS 6.0 kernel loaded in HMA):**

Offset	Size	Description
00h	WORD	offset of next HMA Memory Block (0000h if last block)
02h	WORD	size of this block in bytes (at least 10h)
04h	BYTE	system type (0 = HMA Memory Block as reported by MEMINFO)
05h	BYTE	00h system
06h	BYTE	01h MIB
07h	BYTE	02h NLS/NL
08h	BYTE	03h SHAR1
09h	BYTE	04h LsdMAX
0Ah	BYTE	05h COMMAND
0Bh	BYTE	06h COMMAND
0Ch	BYTE	07h COMMAND
0Dh	BYTE	08h COMMAND
0Eh	BYTE	09h COMMAND
0Fh	BYTE	0Ah COMMAND
10h	BYTE	0Bh COMMAND
11h	BYTE	0Ch COMMAND
12h	BYTE	0Dh COMMAND
13h	BYTE	0Eh COMMAND
14h	BYTE	0Fh COMMAND
15h	BYTE	10h COMMAND
16h	BYTE	11h COMMAND
17h	BYTE	12h COMMAND
18h	BYTE	13h COMMAND
19h	BYTE	14h COMMAND
1Ah	BYTE	15h COMMAND
1Bh	BYTE	16h COMMAND
1Ch	BYTE	17h COMMAND
1Dh	BYTE	18h COMMAND
1Eh	BYTE	19h COMMAND
1Fh	BYTE	1Ah COMMAND
20h	BYTE	1Bh COMMAND
21h	BYTE	1Ch COMMAND
22h	BYTE	1Dh COMMAND
23h	BYTE	1Eh COMMAND
24h	BYTE	1Fh COMMAND
25h	BYTE	20h COMMAND
26h	BYTE	21h COMMAND
27h	BYTE	22h COMMAND
28h	BYTE	23h COMMAND
29h	BYTE	24h COMMAND
2Ah	BYTE	25h COMMAND
2Bh	BYTE	26h COMMAND
2Ch	BYTE	27h COMMAND
2Dh	BYTE	28h COMMAND
2Eh	BYTE	29h COMMAND
2Fh	BYTE	2Ah COMMAND
30h	BYTE	2Bh COMMAND
31h	BYTE	2Ch COMMAND
32h	BYTE	2Dh COMMAND
33h	BYTE	2Eh COMMAND
34h	BYTE	2Fh COMMAND
35h	BYTE	30h COMMAND
36h	BYTE	31h COMMAND
37h	BYTE	32h COMMAND
38h	BYTE	33h COMMAND
39h	BYTE	34h COMMAND
3Ah	BYTE	35h COMMAND
3Bh	BYTE	36h COMMAND
3Ch	BYTE	37h COMMAND
3Dh	BYTE	38h COMMAND
3Eh	BYTE	39h COMMAND
3Fh	BYTE	3Ah COMMAND
40h	BYTE	3Bh COMMAND
41h	BYTE	3Ch COMMAND
42h	BYTE	3Dh COMMAND
43h	BYTE	3Eh COMMAND
44h	BYTE	3Fh COMMAND
45h	BYTE	40h COMMAND
46h	BYTE	41h COMMAND
47h	BYTE	42h COMMAND
48h	BYTE	43h COMMAND
49h	BYTE	44h COMMAND
4Ah	BYTE	45h COMMAND
4Bh	BYTE	46h COMMAND
4Ch	BYTE	47h COMMAND
4Dh	BYTE	48h COMMAND
4Eh	BYTE	49h COMMAND
4Fh	BYTE	4Ah COMMAND
50h	BYTE	4Bh COMMAND
51h	BYTE	4Ch COMMAND
52h	BYTE	4Dh COMMAND
53h	BYTE	4Eh COMMAND
54h	BYTE	4Fh COMMAND
55h	BYTE	50h COMMAND
56h	BYTE	51h COMMAND
57h	BYTE	52h COMMAND
58h	BYTE	53h COMMAND
59h	BYTE	54h COMMAND
5Ah	BYTE	55h COMMAND
5Bh	BYTE	56h COMMAND
5Ch	BYTE	57h COMMAND
5Dh	BYTE	58h COMMAND
5Eh	BYTE	59h COMMAND
5Fh	BYTE	5Ah COMMAND
60h	BYTE	5Bh COMMAND
61h	BYTE	5Ch COMMAND
62h	BYTE	5Dh COMMAND
63h	BYTE	5Eh COMMAND
64h	BYTE	5Fh COMMAND
65h	BYTE	60h COMMAND
66h	BYTE	61h COMMAND
67h	BYTE	62h COMMAND
68h	BYTE	63h COMMAND
69h	BYTE	64h COMMAND
6Ah	BYTE	65h COMMAND
6Bh	BYTE	66h COMMAND
6Ch	BYTE	67h COMMAND
6Dh	BYTE	68h COMMAND
6Eh	BYTE	69h COMMAND
6Fh	BYTE	6Ah COMMAND
70h	BYTE	6Bh COMMAND
71h	BYTE	6Ch COMMAND
72h	BYTE	6Dh COMMAND
73h	BYTE	6Eh COMMAND
74h	BYTE	6Fh COMMAND
75h	BYTE	70h COMMAND
76h	BYTE	71h COMMAND
77h	BYTE	72h COMMAND
78h	BYTE	73h COMMAND
79h	BYTE	74h COMMAND
7Ah	BYTE	75h COMMAND
7Bh	BYTE	76h COMMAND
7Ch	BYTE	77h COMMAND
7Dh	BYTE	78h COMMAND
7Eh	BYTE	79h COMMAND
7Fh	BYTE	7Ah COMMAND
80h	BYTE	7Bh COMMAND
81h	BYTE	7Ch COMMAND
82h	BYTE	7Dh COMMAND
83h	BYTE	7Eh COMMAND
84h	BYTE	7Fh COMMAND
85h	BYTE	80h COMMAND
86h	BYTE	81h COMMAND
87h	BYTE	82h COMMAND
88h	BYTE	83h COMMAND
89h	BYTE	84h COMMAND
8Ah	BYTE	85h COMMAND
8Bh	BYTE	86h COMMAND
8Ch	BYTE	87h COMMAND
8Dh	BYTE	88h COMMAND
8Eh	BYTE	89h COMMAND
8Fh	BYTE	8Ah COMMAND
90h	BYTE	8Bh COMMAND
91h	BYTE	8Ch COMMAND
92h	BYTE	8Dh COMMAND
93h	BYTE	8Eh COMMAND
94h	BYTE	8Fh COMMAND
95h	BYTE	90h COMMAND
96h	BYTE	91h COMMAND
97h	BYTE	92h COMMAND
98h	BYTE	93h COMMAND
99h	BYTE	94h COMMAND
9Ah	BYTE	95h COMMAND
9Bh	BYTE	96h COMMAND
9Ch	BYTE	97h COMMAND
9Dh	BYTE	98h COMMAND
9Eh	BYTE	99h COMMAND
9Fh	BYTE	9Ah COMMAND
A0h	BYTE	9Bh COMMAND
A1h	BYTE	9Ch COMMAND
A2h	BYTE	9Dh COMMAND
A3h	BYTE	9Eh COMMAND
A4h	BYTE	9Fh COMMAND
A5h	BYTE	A0h COMMAND
A6h	BYTE	A1h COMMAND
A7h	BYTE	A2h COMMAND
A8h	BYTE	A3h COMMAND
A9h	BYTE	A4h COMMAND
AAh	BYTE	A5h COMMAND
ABh	BYTE	A6h COMMAND
AC	BYTE	A7h COMMAND
AD	BYTE	A8h COMMAND
AE	BYTE	A9h COMMAND
AF	BYTE	AAh COMMAND
B0	BYTE	ABh COMMAND
B1	BYTE	AC
B2	BYTE	AD
B3	BYTE	AE
B4	BYTE	AF
B5	BYTE	B0
B6	BYTE	B1
B7	BYTE	B2
B8	BYTE	B3
B9	BYTE	B4
BA	BYTE	B5
BB	BYTE	B6
BC	BYTE	B7
BD	BYTE	B8
BE	BYTE	B9
BF	BYTE	BA
C0	BYTE	BB
C1	BYTE	BC
C2	BYTE	BD
C3	BYTE	BE
C4	BYTE	BF
C5	BYTE	C0
C6	BYTE	C1
C7	BYTE	C2
C8	BYTE	C3
C9	BYTE	C4
CA	BYTE	C5
CB	BYTE	C6
CC	BYTE	C7
CD	BYTE	C8
CE	BYTE	C9
CF	BYTE	CA
D0	BYTE	CB
D1	BYTE	CC
D2	BYTE	CD
D3	BYTE	CE
D4	BYTE	CF
D5	BYTE	D0
D6	BYTE	D1
D7	BYTE	D2
D8	BYTE	D3
D9	BYTE	D4
DA	BYTE	D5
DB	BYTE	D6
DC	BYTE	D7
DD	BYTE	D8
DE	BYTE	D9
DF	BYTE	DA
E0	BYTE	DB
E1	BYTE	DC
E2	BYTE	DD
E3	BYTE	DE
E4	BYTE	DF
E5	BYTE	E0
E6	BYTE	E1
E7	BYTE	E2
E8	BYTE	E3
E9	BYTE	E4
EAh	BYTE	E5
ECh	BYTE	E6
ED	BYTE	E7
EE	BYTE	E8
EF	BYTE	E9
F0	BYTE	EAh
F1	BYTE	ECh
F2	BYTE	ED
F3	BYTE	EE
F4	BYTE	EF
F5	BYTE	F0
F6	BYTE	F1
F7	BYTE	F2
F8	BYTE	F3
F9	BYTE	F4
FA	BYTE	F5
FB	BYTE	F6
FC	BYTE	F7
FD	BYTE	F8
FE	BYTE	F9
FF	BYTE	FA

**INT 21h Function 4459h****DR MultiUser DOS 5.0****CP/M compatible API**

This function is used to request the CP/M API provided by Digital Research's MultiUser DOS.

**Call With**

AX 4459h  
 CX 0000h  
 DS:DI parameters

**Notes:** DR-DOS 5.0 does not support the AX=0000h interrupt function.

This API is also available on INT 10h.

**See Also:** 21-4452h

**INT 21h Function 4Ah  
RESIZE MEMORY BLOCK****DOS 2+**

Although undocumented, this function has undocumented behavior when there is not enough memory to satisfy the request.

**Call With**

- AH = 4Ah
- BX = new size in paragraphs
- ES = segment of block to resize

**Returns**

- CF clear if successful
- CF set on error
- AX = error code (07h,08h,09h) (see 21\_59h)
- BX = new program paragraphs available for specified memory block

**Notes:** Under DOS 2.1, through 6.0, if there is sufficient memory to expand the block as much as requested, the block will still be made as large as possible, rather than remaining unchanged.

DOS 2.1.0.0 coalesces any free blocks immediately following the block to be resized.

**See Also:** 21\_40h, 21\_49h

**INT 21h Functions 4B03h and 4B04h  
"EXEC" - LOAD AND/OR EXECUTE PROGRAM****DOS 2+**

In addition to the exec invoked subroutines to execute a child process and load an overlay, the EXEC family also has a routine, undocumented, to load and/or execute an overlaid program without beginning execution. The European OEMs use for 4B03h supports a sequential, unattended subprogram to run a child process in the background.

**Call With**

- AH = 4Bh
- AL = type of load
  - 00h load and execute (documented)
  - 01h load but do not execute
  - 02h load overlay (documented)
  - 03h load and execute (undocumented European MS-DOS 4.0 only)
- ES:BX = program name (see note below) (see Chapter 100/101 for pointer to VCI program name) (16-bit string, 8-bit extension)
- DS:BX = pointer to parameter block (see below)

**Returns**

- CF clear if successful
- BX,DX destroyed
- ES:BX set to 4:1 - process ID set to new program's PSP (see 21\_62h)
- CF set on error
- AX = error code (01h,02h,05h,08h,0Ah,0Eh) (see 21\_59h)

**Notes:** DOS 2.0 supports overlays including SSP.

For a background process, it is possible that there is not enough unallocated memory available, if necessary, by releasing memory with 21/49h or 21/4Ah.

For function 01 - the child process is passed to the child program upon top of the child's stack.

Function 01 has several comments in the early GATP and SSP. Are reserved in the COM structure of DOS 5.0 but was undocumented in previous versions.

For a detailed explanation of function 01b, see Tom Patterson's chapter on the MS-DOS 4.0 manager interface in *Undocumented DOS*, 1st Edition.

NOTE: Arguments such as DR, DCS, etc., check the parameters and parameter block and return an error if an invalid value such as an offset of 4111h is found.

LOAD is a program for E-topics. MS-DOS 4.0 must use the New Executable (NE) format. Note that NE executables are running with the following register values, where the command line is expanded to a valid executable's PSP (0000h) and following, except that the trailing carriage return is turned into a NUL (00h):

AX = environment segment  
 AX = offset of first and last in environment segment  
 AX = offset of automatic data segment (0000h) + 64h  
 BX, BP = 0000h  
 DS = automatic data segment  
 SS, SP = initial stack

**BUG:** DOS 3.0 incorrectly sets DS points at the current program's PSP.

The DOS 3.0 loader does not copy words up to 512 bytes too many if the file contains additional data (0 - 0x10000). The loader's PDATA overlays up to the next 512 byte boundary.

See Also: 21-40h, 21-64h, 21-80h, 1N1-23h

### Format of EXEC parameter block (LOAD) for AL=01h,04h:

Offset	Size	Description
00h	WORD	segment of environment to copy for child process (copy caller's environment at 0000h)
02h	DWORD	pointer to command line to be copied into child's PSP
06h	DWORD	pointer to first 16 bytes to be copied into child's PSP
0Ah	DWORD	pointer to second 16 bytes to be copied into child's PSP
0Eh	DWORD	AX=01h will hold subprogram's initial SS:SP on return
12h	DWORD	AX=04h will hold entry point CS:IP on return

Note: A very subtle change in the EXE/D structure but erroneously swaps the last two fields in the DOS 5.0 *MS-DOS 5.0 Programmer's Reference*. This was corrected in the DOS 6.0 reference.

### Format of .EXE file header:

Offset	Size	Description
00h	2 BYTES	EXE signature, either "MZ" or "ZM" (5A4Dh or 4D5Ah)
02h	WORD	number of bytes in last 512-byte page of executable
04h	WORD	number of 512-byte pages in executable (includes a few pages of system)
06h	WORD	number of relocation entries
08h	WORD	number of paragraphs
0Ah	WORD	number of paragraphs of memory to allocate in addition to executable's size
0Ch	WORD	number of paragraphs of memory to allocate in addition to executable's size
0Eh	WORD	initial SS relative to start of executable
10h	WORD	initial SP
12h	WORD	checksum of executable segment of source file (all words in executable)
14h	WORD	checksum of executable segment of source file (all words in executable)
16h	WORD	checksum of executable segment of source file (all words in executable)
18h	WORD	checksum of executable segment of source file (all words in executable)
1Ah	WORD	40h or greater for new format (NE, FE, PE, W32), executable overlays another's normally (0000h) main program

### —new executable—

1Ch	4 BYTES	unknown
20h	WORD	behavior bits
22h	26 BYTES	reserved for additional behavior/info

3Ch	DWORD	offset of new executable (or other) on MZ header within disk file, or 00000000h if old-style MZ executable
<b>—other—</b>		
1Ch	var	optional information added by linker (program compressor, etc.)
N	DWORDs	relocation items

**Notes.** If the word at offset 02h is 4, it should be treated as 00h since pre-1.10 versions of the MS linker set it that way.

If both the minimum and maximum allocation (offset 0Ah/0Ch) are zero, the program is linked as high as memory as possible. The Windows KERNL uses this; see Pietrek, *Windows Internals*.

The maximum allocation is set to FFFh by default.

The new executable (NE) and COFF executable (COFF) headers are described in detail in the on-disk interrupt listing. The NE and Win32 portable executable (PE) formats are both documented by Microsoft.

## INT 21h Function 4Eh

**DOS 2+**

### "FINDFIRST" FIND FIRST MATCHING FILE

Although often treated as a function, it has undocumented fields in its data structure, which are used to record the progress of the directory search; an undocumented `error` and `errorflag`.

#### Call With

- AX: 4Eh
- AX: special flag (see the APPENDIX note below)
- CX: file attribute mask (bits 0 and 5 cleared)
- DS:DX: pointer to AMIZ file specification; may include path and wildcard

#### Returns

- CF: clear if successful
- DI:EA: findfirst data block (see below)
- CF: set on error
- AX: error code (02h, 03h, 12h) (see 2F/50h)

**Notes.** For search attributes, bits that double all files with a MUI file specification (subroutine of hidden system) and directory attributes will be returned. Unlike DOS 2.x searching for attribute 08h (archive), it will also return normal files when under DOS 3.x with the attribute set to 0x.

This call also returns successfully if given the name of a character device without wildcards. DOS 2.x returns attribute 00h, size 0, and the current date and time. DOS 3.x returns attribute 40h and the current date and time.

Immediate data at 2F/50h of APPENDIX returns the same; the same as DS:DX and on overflow, a 0x00000000. The actual found path and will be stored otherwise. The actual found path is 16 bytes per file to the original space without a path.

Under VNT, it may be used to generate a server-side test results by searching for "SERVER\*" as a string for resources may be obtained. The searching for "SERVER\*" is "BUG" (see DOS 3.x and 4.x) to search files frequently. It is this list or with a character device (e.g., "0x00000000" search file by which you can do the same case (e.g., 08h) will find unless it is a character device (DOS call which specifies an open performance directory search with out the volume file "0x00000000" searches are performed by CTRL+ALT+2F/50h OPEN 2F/50h UNLINK 2x/50h and RENAME 2F/50h.

**See Also:** 2F/11h, 2F/AFh, 2F/111Bh, 2F/B711h

**Format of FindFirst data block (in DTA):**

Offset	Size	Description
<b>—PCDOS 3.10, PCDOS 4.01, MS DOS 3.2/3.3/5.0—</b>		
00h	BYTE	drive letter - bits 6-0s: remove if bit 7 set
01h	11 BYTES	search template
0C h	BYTE	search attributes
<b>—DOS 2.x (and some DOS 3.x)—</b>		
00h	BYTE	search attributes
01h	BYTE	drive letter
02h	11 BYTES	search template
<b>—DOS 2.x and most 3.x—</b>		
01h	WORD	entry count within directory
03h	DWORD	pointer to DTA
05h	WORD	entry count of root directory
<b>—PCDOS 4.01, MS DOS 3.2/3.3/5.0—</b>		
01h	WORD	entry count within directory
01h	WORD	entry count of root directory
01h	BYTE	entry count
<b>all versions documented fields</b>		
05h	BYTE	attribute of file found
10h	WORD	file size bits 11-15: hour bits 5-10: minute bits 0-4: seconds / 2
18h	WORD	file date bits 9-15: year 1980 bits 5-8: month bits 0-4: day
1Ah	DWORD	file
1Ch	13 BYTES	ASCIZ filename-extension

**INT 21h Function 50h****DOS 2+****SET CURRENT PROCESS ID (SET PSP ADDRESS)**

Force a new value for DOS's current process ID (the current process's PSP segment), thus effectively becoming another process.

**Call With**

AX	50h
BX	segment of PSP for new process

**Notes:** DOS sets the current process ID to the value of the current process's PSP segment, thus effectively becoming another process. This function is used by other DOS programs to change the current process ID without setting the current process ID to the value of the current process's PSP segment.

**Critical Error:** If the current process ID is not set to the value of the current process's PSP segment, the current process ID is not set to the value of the current process's PSP segment.

at any time, even during a normal INT 21h call. See Chapter 6, Figure 6-4.)

Some MS-DOS versions (3.21, 3.31, 5.0) do not support this function. If you are using one of these versions, you should only call this function if you are running a program that is designed to handle invalid addresses. See PSP/FAT/C in Chapter 4.

Some versions of DOS (3.21, 3.31, 5.0) do not support this function. If you are using one of these versions, you should only call this function if you are running a program that is designed to handle invalid addresses. See PSP/FAT/C in Chapter 4.

Some versions of DOS (3.21, 3.31, 5.0) do not support this function. If you are using one of these versions, you should only call this function if you are running a program that is designed to handle invalid addresses. See PSP/FAT/C in Chapter 4.



This call was undocumented prior to the release of DOS 6.0.

See Also: 2, 26x 21, 51h, 21, 65h

---

## INT 21h Function 51h DOS 2+ GET CURRENT PROCESS ID (GET PSP ADDRESS)

Returns the segment address of the current process's PSP, which is used by DOS as a process identifier.

Call With

AH 51h

Returns:

ES segment of PSP for current process.

**Notes:** DOS uses the current PSP address to determine which processes own files and folders. It corresponds to process identifiers used by other OSs.

Under DOS 2.x, this function cannot be invoked inside an INT 2Bh handler without setting a Critical Error flag.

Under DOS 3.x, this function does not use any of the DOS internal stacks and may thus be called at any time, even during another INT 21h.

This function is identical to the documented function 21, 62h, and is supported by OS/2 compatibility box.

This call was not documented for DOS 2.x/4.x, but has been available since DOS 5.0.

See Also: 2, 26x 21, 50h, 21, 62h

---

## INT 21h Function 52h DOS 2+ "SYSVARS" - GET LIST OF LISTS

Determines the address of DOS's internal list of tables and lists. Most internal list structures are reachable through this list.

Call With

AH 52h

Returns:

ES:BX pointer to DOS list of lists.

**Note:** This function is partially supported on some OS/2 1.11 compatibility systems. However, on some platforms, such as EMM386, MSDOSM is not available, and the CPU kernel may not be EMM386. The implementation of this structure is simple; see Figure 1.9.

### Format of List of Lists (SysVars):

Offset	Size	Description
12	WORD	DOS 5.0's 80-starting read count; see 71, 4000h.
10	WORD	DOS 3.x's 80-starting read count; see 71, 4000h.
8	DWORD	DOS 5.x's 80-starting read count.
4	WORD	DOS 3.x's 80-starting read count.
		When CON is read via a handle, DOS reads an entire line and returns the respective portion. EOF is returned for the next read. 0000h indicates no internal input segment of first or more control blocks.
2	WORD	pointer to first DOS filename block; see INT 21, AH 32h.
00h	DWORD	pointer to first Sector File; see below.
04h	DWORD	pointer to active CON device segment of second CON object.
08h	DWORD	drive with CLOCK bit (bit 5, set).
0Ch	DWORD	pointer to active CON device segment of second CON object with STDIN bit (bit 0, set), such as ANSISYS.

**DOS 2.x--**

60	BYTE	number of logical drives in system
1	WORD	maximum bytes / block of any block device
5	DWORD	pointer to first disk buffer (see below)
	STRUCT	pointer to system FC B tables (not a pointer!) NLT is always the first device on DOS's linked list of device drivers (see below for format)

**—DOS 3.0—**

10	BYTE	number of block devices
16	WORD	maximum bytes / block of any block device
5	DWORD	pointer to first disk buffer (see below)
7	STRUCT	pointer to system FC B tables (see below)
1	BYTE	value of <code>FASTDRIVE</code> command in <code>CONFIG.SYS</code> (default 5)
0	DWORD	pointer to <code>STRING</code> workspace area
60	WORD	value of <code>STRING</code> (to <code>STRING</code> from <code>CONFIG.SYS</code> )
55	DWORD	pointer to FC B table
0	WORD	the <code>y</code> in <code>FCBS-y</code> from <code>CONFIG.SYS</code>
58	STRUCT	pointer to NLT device driver header (not a pointer!) NLT is always the first device on DOS's linked list of device drivers (see below for format)

**DOS 3.1 3.3--**

10	WORD	maximum bytes per sector of any block device
170	DWORD	pointer to first disk buffer (not a pointer) (see below)
102	DWORD	pointer to system FC B tables (not a pointer) (see below)
130	DWORD	pointer to system FC B tables (see below)
0	WORD	the <code>y</code> in <code>FCBS-y</code> from <code>CONFIG.SYS</code>
70	BYTE	number of block devices installed
21	LIST	list of system FC B tables (up to 5 installed block devices and <code>CONFIG.SYS</code> <code>FASTDRIVE</code> ). Also size of current directory structure array
27	IRBYTES	actual NLT device driver header (not a pointer!) NLT is always the first device on DOS's linked list of device drivers (see below for format)
34	BYTE	number of JOINed drives

**DOS 4.x**

100	WORD	maximum bytes per sector of any block device
120	DWORD	pointer to disk buffer (into record) (see below)
160	STRUCT	pointer to system FC B tables (not a pointer) (see below)
151	DWORD	pointer to system FC B tables (see below)
110	LIST	list of system FC B tables (up to 5 installed FCBS-y's (always 00h for DOS 3.0)
200	BYTE	number of block devices installed
210	LIST	list of system FC B tables (up to 5 installed block devices and <code>CONFIG.SYS</code> <code>FASTDRIVE</code> ). Also size of current directory structure array
720	STRUCT	pointer to NLT device driver header (not a pointer!) NLT is always the first device on DOS's linked list of device drivers (see below for format)
340	BYTE	number of JOINed drives

35h	WORD	pointer within MS-DOS SYS data segment to list of special program names (see below); always 0000h for DOS 5.0-6.0.
37h	DWORD	pointer to FAR output file inside of IPS if its function (see below). This row may be called by any IPS driver which does not wish to service functions 20h or 24h. 28h itself.
38h	DWORD	pointer to chain of IPS available to system drivers.
39h	WORD	the value of FIFRS via rounded up to multiple of 30 (for EMS).
41h	WORD	vector of overlaid buffers; the value of FIFRS via and offsets 38h and 39h under "DOS 5.0-6.0" should read the value of FIFRS via.
43h	BYTE	boot drive (A-Z).
44h	BYTE	00h if 80386+; 00h otherwise.
45h	WORD	extended memory size in K.

## —DOS 5.0-6.0—

10h	39 BYTES	as for DOS 4.x (see above).
37h	DWORD	pointer to SEVER program; 38 or 0000h/0000h.
38h	WORD	DOS FIFRS offset to DOS CS of function to fix A20 control when executing special COM format.
39h	WORD	PSP of most recent EXEC'd program if DOS = HIMEM; 0000h if DOS low.
3Fh	8 BYTES	as for DOS 4.x (see above).

**Format of memory control block (see also below):**

Offset	Size	Description
00h	BYTE	block type; 5Ah if last block in chain; otherwise 41h.
01h	WORD	PSP segment of owner or 0000h if free.
		0000h if DR-DOS XMS UMB.
		0001h if DR-DOS exclusive upper memory "swap".
		0008h if belongs to DOS.
		111Ah if 386MAX UMB control block.
		111Dh if 386MAX locked out memory.
		111Eh if 386MAX UMB reserved area follows by control block.
03h	WORD	size of memory block in paragraphs.
05h	3 BYTES	unused.

## —DOS 2.x, 3.x—

00h	8 BYTES	unused.
-----	---------	---------

## —DOS 4.x—

00h	8 BYTES	ASCII program name; if PSP memory block or DR-DOS UMB, it is garbage null terminated if less than 8 characters.
-----	---------	---

**Notes:** The next ACh is a segment, current + size + 1.

Under DOS 3.1-3.x, first memory blocks in the DOS data segment containing nonalphanumeric bytes etc. Under DOS 4.x it is divided into subsegments, each with its own memory control block (see below); the first of which is at offset 0000h.

For DOS 5.0 blocks owned by DOS may have either "SD" or "SD" in bytes 08+ and 09h. "SD" is system code, or locked out under UMB memory. "SD" is system data device driver, etc.

Some versions of DR-DOS use only seven characters of the program name, placing a NUL in the eighth byte.

**Format of MS-DOS 5.0 UMB control block:**

Offset	Size	Description
00h	WORD	type - SM if last block in chain, 4Dh otherwise
01h	WORD	first available paragraph in UMB if control block at start of UMB, 000Ah if control block at end of UMB
03h	WORD	length in paragraphs of following UMB or wicket-out region
05h	2 BYTES	unused
06h	8 BYTES	block type name "UMB" if start block "SM" if end block in UMB

**Format of DOS 4+ data segment subsegment control blocks:**

Offset	Size	Description
00h	WORD	subsegment type - blocks typically appear in this order
01h	WORD	"D" device driver
02h	WORD	"E" device driver appendage
03h	WORD	"I" IFS (installable file system) driver
04h	WORD	"F" FLS - control block storage area for FIFS-5
05h	WORD	"X" XCRS - control block storage area, if present
06h	WORD	"C" CUFFERS - workspace area of BUFFERS.A (option used)
07h	WORD	"B" BUFFERS - storage area
08h	WORD	"D" DIRECTORY - control of directory structure array storage area
09h	WORD	"S" STACKS - code and data area, if present (see below)
0Ah	WORD	"T" INSTALL - transient code
0Bh	WORD	paragraph 05 (segment start) usually the text paragraph
0Ch	WORD	size of subsegment in paragraphs
0Eh	2 BYTES	unused
0Fh	8 BYTES	for names "D" and "E" - base name of file from which the driver was loaded (unused for other types)

**Format of data at start of STACKS code segment (if present):**

Offset	Size	Description
00h	WORD	unknown
02h	WORD	number of stacks (the <i>x</i> in STACKS= <i>x</i> )
04h	WORD	<i>x</i> of stack control block array, should be 8*x
06h	WORD	size of each stack (the <i>y</i> in STACKS= <i>x</i> )
08h	DWORD	pointer to STACKS data segment
0Ch	WORD	offset to STACKS data segment of stack control block array
0Eh	WORD	offset to STACKS data segment of first element of that array
10h	WORD	offset to STACKS data segment of the control array for the next stack to be allocated (initially same as value 00Ch, and works as a window in steps of 16 to the value 000Ch in a 64K word memory, preempts each other)

Note 1 - STACKS = *x* - control block array, if present, is located as follows:

- DOS 3.2** - For subsegment *code* (text paragraph) base address *base* in the IBMPC segment (seen at 00700190h)
- DOS 3.3** - For code segment *code* (text paragraph) base address *base* in DOS data segment which must be determined by inspecting the segment pointers of the vectors for those of type "P" (01, 02, 03, 04, 05, 06, 07, 08, 09, 0A, 0B, 0C, 0D, 0E, 0F) which have not been redirected by device drivers or ISRs.
- DOS 4.x** - Identified by subsegment *code* (text paragraph) type "S" within the DOS data segment

**Format of array elements in STACKS data segment:**

Offset	Size	Description
00h	BYTE	status: 00h=free, 01h=in use, 03h= corrupted by overflow of a higher stack.
01h	BYTE	not used
02h	WORD	previous SP
04h	WORD	previous SS
06h	WORD	ptr to word at top of stack, new value for SP's 16x word at the top of the stack is preset to point back to this control block.

**SHARE.EXE hooks (DOS 3.1-6.0):**

offsets from first System File Table pointer at file: \system\sys-04h; see SHAREHOOK.C in Chapter 8.

Offset	Size	Description
7Ch	DWORD	pointer to FAR routine for <i>unlock a p/p file</i> Note: not called by MS-DOS 3.3; set to 00000h:00000h by SHARE 3.3
38h	DWORD	pointer to FAR routine called on opening file on call internal DOS location points at <i>tkram</i> ; see 21-5106h <b>Returns:</b> CF clear if successful CF set on error AX: DOS error code: 24h (see 21-506h) Note: SHARE assumes DS:SS:DI:SI and directly accesses DOS internally to get the name of the file requested.
34h	DWORD	pointer to FAR routine called on closing file FS:DI pointer to system file table Note: SHARE does something to <i>unlock</i> record for file.
40h	DWORD	pointer to FAR routine to close all files for <i>gave</i> computer called by 21-5103h.
2Ch	DWORD	pointer to FAR routine to close all files for <i>gave</i> process called by 21-5104.
28h	DWORD	pointer to FAR routine to close file <i>work</i> called by 21-5102. DS:SI pointer to DOS parameter list; see 21-5100h DI+DI*DX pointer to name of file to close <b>Returns:</b> CF clear if successful CF set on error AX: DOS error code (03h); see 21-509
24h	DWORD call with BX CX:DX SI:AX	pointer to FAR routine to lock region of local file file handle starting offset size <b>Returns:</b> CF set on error AX: DOS error code: 21h; see 21-506 Note: not called if file is marked as remote
20h	DWORD call with BX	pointer to FAR routine to unlock region of local file file handle

	CX:DX SI:AX	starting offset size <b>Returns:</b> CF set on error AX - DOS error code (21h) (see 21-50h)
1C h	<b>DWORD</b> call with DS:DI CX	<b>Note:</b> not called by current system (see 21-110B) pointer to FAR routine to check if file is locked pointer to system file table entry for file length of region from current position in file <b>Returns:</b> CF set if any portion of region locked AX - 0021h
1D h	<b>DWORD</b>	pointer to FAR routine to get open file list entry called by 21-5100h DS:SI pointer to DOS parameter list (see 21-5100h) DI % BX index of sharing record DI % CX index of SE1 in SE1 chain of sharing rec. <b>Returns:</b> CF set on error or not loaded AX - DOS error code (12h) (see 21-50h) CF clear if successful FS:DI pointer to filename CX - number of locks owned by specified SE1 BX - network machine number DX - destroyed
14 h	<b>DWORD</b>	pointer to FAR routine for updating I-C-B from SE1 call with DS:SI pointer to unopened I-C-B FS:DI pointer to system file table entry <b>Returns:</b> BF = 0h <b>Note:</b> copies following fields from SE1 to I-C-B starting cluster of file 00h 14h sharing record offset 33h 1Ch file attribute 04h 11h
10 h	<b>DWORD</b>	pointer to FAR routine to get first cluster of I-C-B file call with FS:DI pointer to system file table entry DS:SI pointer to I-C-B <b>Returns:</b> CF set if SE1 cluster or sharing record offsets mismatched CF clear if successful starting cluster number from I-C-B
0C h	<b>DWORD</b> DS:SI	FS:DI pointer to system file table <b>Returns:</b> AX - number of handles in IFH which already uses SE1 <b>Note:</b> called during open, create of a file <b>Note:</b> FS:SI is not updated with file name embedded in system file table. SE1 table does something to a file or not. SE1 table does not know what type the same file open mode and sharing record
02 h	<b>DWORD</b>	pointer to unknown FAR routine

04h	DWORD	Note: closes various handles (referring to file, most recently opened pointer to FAR routine to update directory and via related SEI entries) call with FSI DI pointer to system file table entry for file (see below) AX = subfunction (apply to each related SEI) 00h: update time stamp (offset 01E) and date stamp (offset 01h) 01h: update file size (offset 113) and start sector (offset 0B5) Sets 1st accessed cluster to ds to start of file (file never accessed) 02h: as function 01h, but last accessed holds a way's char. gen. 03h: do both functions 00h and 02h Note: follows pointer at offset 7B5 in system file table entries Note: NOT if opened with no inherit or via FCB
-----	-------	--

Note: Most of the above books or SS-DOS-DOS DS, and directly except 04h, 14h, 18h and 30h assume either that SS-DOS DS access DOS internal data

#### Format of sharing record:

Offset	Size	Description
00h	BYTE	flag 00h free block 01h allocated block 1Fh end marker
01h	WORD	size of block
03h	BYTE	checksum of path name (including NUL) if sum of ASCII values is N, checksum is $N * 256 + N * 256$
04h	WORD	offset in SHARE.SDS of lock record (see below)
06h	DWORD	pointer to start of system file table chain for file
0Ah	WORD	unique sequence number
0C h	var	ASCII file path name

#### Format of SHARE.EXE lock record:

Offset	Size	Description
00h	WORD	offset to SHARE.SDS of next lock table in use
02h	DWORD	offset to file of start of locked region
06h	DWORD	offset to file of end of locked region
0Ah	DWORD	pointer to system file table entry for this file
0E h	WORD	FSP segment of lock owner

#### Format of DOS 2.x system file tables:

Offset	Size	Description
00h	WORD	pointer to next file table (offset + FFFFh if last)
04h	WORD	number of files in this table
06h	28h	bytes per file
Offset	Size	Description
00h	BYTE	number of file handles referring to this file
01h	BYTE	file open mode (see 2F-30h)
02h	BYTE	file attribute
03h	BYTE	drive 0 (abstract device) 1 A 2 B etc.
04h	11 BYTES	filename in FCB format (no path, no period, blank padded)
0Fh	WORD	unknown
11h	WORD	unknown
13h	DWORD	file size
17h	WORD	file data in packed format (see 2F-5-00h)

10h	WORD	file time in packed format (see 2f 5700h)
11h	BYTE	device attribute (see 2f 4400h)
<b>character device—</b>		
14h	DWORD	pointer to device driver
<b>—block device—</b>		
17h	WORD	starting cluster of file
18h	WORD	relative cluster in file of last cluster accessed
20h	WORD	absolute cluster number of current cluster
22h	WORD	<i>unknown</i>
24h	DWORD	current file position

#### Format of DOS 3.0 system file tables and FCB tables

Offset	Size	Description
00h	DWORD	pointer to next file table; offset 1111h if last
04h	WORD	number of files in this table
06h		38h bytes per file
Offset	Size	Description
00h	11h as of DOS 3.1+; see below	
11h	WORD	byte offset of directory entry within sector
14h	BYTE	file attribute; bit 0 set means packed file
24h	DWORD	SHARE FCB pointer; same as same file
30h	WORD	file size; see 2f 1600h; use the virtual file size if the file is a stream file; see 2f 1600h
32h	WORD	file creation time; see 2f 5700h
34h	WORD	file last access time; see 2f 5700h; apparently always 0000h

#### Format of DOS 3.1 3x system file tables and FCB tables

Offset	Size	Description
0h	WORD	pointer to next file table; offset 1111h if last
04h	WORD	number of files in this table
06h		35h bytes per file
Offset	Size	Description
00h	WORD	number of file handles referring to this file
02h	WORD	file open mode (see 2f 3f00h)
04h	BYTE	bit 15 set if this file opened via FCB
05h	WORD	file attribute
06h	WORD	device info word (see 2f 4400h)
08h	WORD	bit 18 set if remote file
09h	WORD	bit 14 set means do not set file date/time on closing
0ah	WORD	bit 12 set means don't inherit on EXEC
0ch	WORD	bit 8-0 drive number for disk files
0eh	WORD	file pointer to DOS Drive Parameter Block (see 2f 32h)
08h	WORD	starting cluster of file
010h	WORD	file time in packed format (see 2f 5700h)
011h	WORD	file date in packed format (see 2f 5700h)
012h	WORD	file size



**—system file table—**

15h      DWORD      current offset in hlk. may be larger than size of file  
 21-42h does not check new position!

**—FCB table—**

15h      WORD        counter for last L/O to FCB

17h      WORD        counter for last open of FCB

(these are separate to determine the times of the latest I/O and open)

19h      WORD        relative cluster within file of last cluster accessed

1Bh      WORD        absolute cluster number of last cluster accessed  
*0000h if file never read or written*

1Dh      WORD        number of sector containing directory entry

1Eh      BYTE        number of directory entries per sector (set offset 32)

20h      11 BYTE     file name in FCB format (no path, no root, blank padded)

21h      DWORD       SHARE FCB pointer to previous FCB starting same way

2Fs      WORD        SHARE FCB network attach number which opened file

(Windows Enhanced mode/DOSMMGR uses the virtual machine ID as the attach number, see 21-168h)

31h      WORD        PSP segment of file's owner, see 21-26h

(first three entries for AUX/CON/PRN contain segment of DOS system startup code, sometimes contains PSP's MCB rather than PSP itself)

43h      WORD        offset within SHARE FCB of segment of startup record, see above. 000h if none

**Format of DOS 4.0-6.0 system file tables and FCB tables:**

Offset	Size	Description
00h	DWORD	points to next file table, offset = FFFFh if last
04h	WORD	number of files in this table
06h	3Bh bytes per file	
Offset	Size	Description
00h	WORD	number of file handles referring to this file
02h	WORD	file open mode, see 21-170h
04h	BYTE	bit 4 set if this file opened via FCB
05h	WORD	file attributes
		device id (word), see 21-4400h
		bit 15 set if remote file
		bit 14 set if readonly, not set if file data time encoding
		bit 13 set if named pipe
		bit 12 set if no inherit
		bit 11 set if network spooler
		bits 5-0 drive number for disk files
07h	DWORD	pointer to device driver header of character device else pointer to DOS Drive Parameter Block, see 21-322h or RFDIR data
08h	WORD	start sector of file
0Dh	WORD	file time in packed format, see 21-3700h
0Fh	WORD	file date in packed format, see 21-3700h
11h	DWORD	file size
15h	DWORD	current offset in hlk

<b>local file—</b>		
19h	WORD	relative cluster within file of last cluster accessed
1Bh	DWORD	number of sector containing directory entry
1Fh	BYTE	number of dir entry within sector (byte offset/32)
<b>—network redirector—</b>		
19h	DWORD	pointer to RE DIRHS record
1Eh	3 BYTES	<i>unknown</i>
20h	11 BYTES	character string in format: no path period blank padded
28h	DWORD	SHARE FILE pointer to previous SMI sharing same file
2Ch	WORD	SHARE FILE network structure number which opened file (Windows Enhanced mode DOSMGR uses the same structure ID as the machine number; see 2F-1088) PSP segment of file's owner (see 21/26h)
31h	WORD	first three entries for ATTR, CDS, PRN contain segment of IO SYS startup code)
35h	WORD	offset within SHARE FILE code segment of sharing record (see above) 0000h if none
39h	WORD	local absolute cluster number of last cluster accessed <i>redirector: unknown</i>
3Fh	DWORD	pointer to HS driver for file 00000000h if native DOS

### Format of current directory structure (CDS) (array, LASTDRIVE entries):

Offset	Size	Description											
00h	67 BYTES	ANALOGUE TO DIRVX PATH local or MGR PATH network drive attributes (see also note below and 21-5f07h)											
43h	WORD	<table border="0"> <tr> <td> <table border="0"> <tr> <td>1</td> <td>physical drive</td> <td rowspan="2">file system of 11</td> </tr> <tr> <td>15</td> <td>JOINS.ct</td> </tr> </table> </td> <td> <table border="0"> <tr> <td>1</td> <td>added if not under MS-DOS or JOIN</td> </tr> <tr> <td>7</td> <td>redirector, but not network (e.g. CD-ROM)</td> </tr> </table> </td> </tr> </table>	<table border="0"> <tr> <td>1</td> <td>physical drive</td> <td rowspan="2">file system of 11</td> </tr> <tr> <td>15</td> <td>JOINS.ct</td> </tr> </table>	1	physical drive	file system of 11	15	JOINS.ct	<table border="0"> <tr> <td>1</td> <td>added if not under MS-DOS or JOIN</td> </tr> <tr> <td>7</td> <td>redirector, but not network (e.g. CD-ROM)</td> </tr> </table>	1	added if not under MS-DOS or JOIN	7	redirector, but not network (e.g. CD-ROM)
<table border="0"> <tr> <td>1</td> <td>physical drive</td> <td rowspan="2">file system of 11</td> </tr> <tr> <td>15</td> <td>JOINS.ct</td> </tr> </table>	1	physical drive	file system of 11	15	JOINS.ct		<table border="0"> <tr> <td>1</td> <td>added if not under MS-DOS or JOIN</td> </tr> <tr> <td>7</td> <td>redirector, but not network (e.g. CD-ROM)</td> </tr> </table>	1	added if not under MS-DOS or JOIN	7	redirector, but not network (e.g. CD-ROM)		
1	physical drive	file system of 11											
15	JOINS.ct												
1	added if not under MS-DOS or JOIN												
7	redirector, but not network (e.g. CD-ROM)												
45h	DWORD	pointer to Drive Parameter Block for drive (see 2/32h)											
<b>—local drives—</b>													
49h	WORD	starting cluster of current directory 0000h if root 1111h if never accessed											
4Bh	WORD	<i>apparently always FFFFh</i>											
4Fh	WORD	<i>apparently always FFFFh</i>											
<b>network drives—</b>													
49h	DWORD	pointer to RE DIRHS record or FFFFh FFFFh <i>stored user data from 21-5F07h</i>											
4Bh	WORD												
4Fh	WORD	<table border="0"> <tr> <td>1</td> <td>path separator, i.e., path of backslash corresponding to root directory for drive</td> </tr> <tr> <td>15</td> <td>number of characters to hide from the C:\DIR \w C:\DIR calls normally set to 2 to hide the 2 characters of the S:\\$ JOIN and networks call paths. C:\ calls appropriate portion of the true path is visible to the user</td> </tr> </table>	1	path separator, i.e., path of backslash corresponding to root directory for drive	15	number of characters to hide from the C:\DIR \w C:\DIR calls normally set to 2 to hide the 2 characters of the S:\\$ JOIN and networks call paths. C:\ calls appropriate portion of the true path is visible to the user							
1	path separator, i.e., path of backslash corresponding to root directory for drive												
15	number of characters to hide from the C:\DIR \w C:\DIR calls normally set to 2 to hide the 2 characters of the S:\\$ JOIN and networks call paths. C:\ calls appropriate portion of the true path is visible to the user												
<b>—DOS 4+—</b>													
51h	BYTE	<i>unknown, used by network</i>											
52h	DWORD	pointer to HS driver of this drive 00000000h if native DOS											
56h	WORD	<i>unknown</i>											

**Notes:** The switch for unabled drives is normally set to `X` (but may be empty after `JOIN x /D` in DR-DOS 5.0 or `NFT USE x /D` in older LAN versions).

Normally, only one of bits 13&12 may be set together with bit 14, but DR-DOS 5.0 uses other combinations (see below).

#### Format of DR-DOS 5.0-6.0 current directory structure entry (array):

Offset	Size	Description
00h	6 <sup>7</sup> BYTES	ASCII path name of actual root directory for this logical drive
42h	WORD	drive attributes 1000h: SL (S)Fcd drive 4000h: JOINed drive 4000h: physical drive 5000h: ASSIGNed drive 7000h: JOINed drive 8000h: network drive
45h	BYTE	physical drive number (0- <code>X</code> ) if this logical drive is valid
46h	BYTE	appears only due to <code>JOIN</code> and <code>ASSIGN</code>
47h	WORD	cluster number of start of parent directory (0000h if root)
49h	WORD	entry number of current directory in parent directory
4Bh	WORD	cluster number of start of current directory
4Dh	WORD	used for media change detection
4Fh	WORD	cluster number of SL (S)Fcd JOINed root directory 0000h if physical root directory

#### Format of device driver header:

Offset	Size	Description
00h	WORD	points to next driver (offset 1FFF of last driver)
04h	WORD	device attributes Character device bit 15: set bit 14: IOCTL supported (see 21, 44h) bit 13: (DOS 3+) output until busy supported bit 12: reserved bit 11: (DOS 3+) (OPEN, CLOSE) Rec/Media calls supported bits 10:8: reserved bit 7: (DOS 5+) (GENERIC_IOCTL) check call supported command 19h bit 6: (DOS 3.2+) Generic_IOCTL call supported command 13h bit 5: reserved bit 4: device supports IOCTL 29h "last console output" bit 3: device is CIO (KS) all reads/writes use transfer record described below bit 2: device is SSI bit 1: device is standard output bit 0: device is standard input Block device bit 15: clear bit 14: IOCTL supported bit 13: non-IBM format bit 12: reserved bit 11: (DOS 3-) (OPEN, CLOSE) Rec/Media calls supported

bit 10 reserved  
 bit 9 *drive 1 (C) not allowed*  
     (set by DOS 3.3 DRIVER SVS for "new" drives)  
 bit 8 *unknown set by DOS 3.3 DRIVER SVS for "new driver*  
 bit 7 DOS 5+ *Generic IOK 11 check call supported*  
     command 19h  
 bit 6 DOS 3.2+ *Generic IOK 11 call supported*  
     command 13h  
     implies support for commands 17h and 18h  
 bits 5-2 reserved  
 bit 1 DOS 3.31+ *drive supports 32-bit sector addressing*  
 bit 0 reserved

06h WORD device strategy entry point  
 08h WORD device interrupt entry point  
 + 0x015EA pointer to request header (see 2F/0802h)

#### —character device—

0A BYTE *device name* disk-padded character device name

#### —block device—

0A BYTE *units* number of subunits drives supported by driver  
 0Bh BYTE *unused*

2h WORD CDDROM driver reserved must be 0000h  
 4h BYTE CDDROM driver *drive 816h* must initially be 00h  
 15h BYTE (CDDROM driver) number of units  
 16h BYTE CDDROM *track signature MSB* Data where 'm' is version (currently '00')

#### Format of CLOCKS transfer record:

Offset	Size	Description
00h	WORD	number of days since 1 Jan 1980
02h	BYTE	minutes
03h	BYTE	hours
04h	BYTE	hundredths of second
05h	BYTE	seconds

#### Format of DOS 2.x disk buffer:

Offset	Size	Description
00h	DWORD	pointer to next disk buffer (offset = FFFFh if last least-recently-used buffer is first in chain)
04h	BYTE	drive (0=A, 1=B, etc., FFFh if not in use)
05h	WORD	<i>applicable to buffered sector always to be 00h/00h/01h</i>
06h	WORD	logical sector number
07h	BYTE	number of copies to write (1 for non-EAT sectors)
08h	BYTE	sector offset between copies if multiple copies to be written
0Ch	DWORD	pointer to DOS Drive Parameter Block (see 21-32h)
10h		buffered data

#### Format of DOS 3.x disk buffer:

Offset	Size	Description
00h	DWORD	pointer to next disk buffer (offset = FFFFh if last least-recently-used buffer is first in chain)
04h	BYTE	drive (0=A, 1=B, etc., FFFh if not in use)

05h	BYTE	flags <i>bit 7 unknown</i> bit 6: buffer dirty bit 5: buffer has been referenced <i>bit 4 unknown</i> bit 3: sector in data area bit 2: sector in a directory, either root or subdirectory bit 1: sector in FAT <i>bit 0: boot sector</i>
06h	WORD	logical sector number
08h	BYTE	number of copies to write, 1 for non FMS sectors
09h	BYTE	sector offset, between copies of multiple copies mode, written pointer to DOS Drive Parameter Block, <i>sector offset = 32h</i>
0A5h	DWORD	pointer to DOS Drive Parameter Block, <i>sector offset = 32h</i>
0EB	WORD	<i>apparently unused, a zero always</i>
10h		buffered data

**Format of DOS 4.00 (pre UR 25066) disk buffer info:**

Offset	Size	Description
00h	DWORD	pointer to address of disk buffer hash chains, <i>reads sector below</i>
04h	WORD	number of disk buffer hash chains, <i>referred to as NDIR_H below</i>
06h	DWORD	pointer to look ahead buffer, zero if not present
0A0	WORD	number of look ahead sectors, also referred to as BU_LHRN_A
0C0h	BYTE	00h if buffers in FMS, <i>sector offset</i>
0D0h	WORD	FMS number for buffers, zero if not in FMS
0Fh	WORD	FMS physical page number used for buffers, usually 255
1Eh	WORD	<i>apparently always 0000h</i>
17h	WORD	segment of FMS physical page frame
18h	WORD	<i>apparently always zero</i>
170	4 WORDS	FMS physical page mapping information

**Format of DOS 4.01****(from UR 25066 Corrective Services Disk on) disk buffer info:**

Offset	Size	Description
00h	DWORD	pointer to address of disk buffer hash chains, <i>reads sector below</i>
04h	WORD	number of disk buffer hash chains, <i>referred to as NDIR_H below</i>
06h	DWORD	pointer to look ahead buffer, zero if not present
0A0h	WORD	number of look ahead sectors, also referred to as BU_LHRN_A
0C0h	BYTE	00h possibly non-integer, from pre-UR 25066 format
0D0	WORD	FMS number for buffers, <i>sector offset</i>
0Fh	WORD	FMS physical page number, <i>FMS number, subsector</i>
1Eh	WORD	FMS number for buffers, <i>sector offset</i>
170	WORD	FMS physical page number, <i>sector offset</i>
18h	WORD	segment of FMS physical page frame
19h	WORD	segment of one sector of workspace buffer, <i>allocated to page memory if BU_LHRN/XS or /XD options in effect, possibly to avoid DMA into FMS</i>
1Ah	WORD	FMS physical page number used for buffers, usually 255
1C0	WORD	<i>appears always to be 0000h</i>
1Eh	WORD	segment of FMS physical page frame
1F0	WORD	<i>appears always to be zero</i>
22h	BYTE	00h if /XS, 01h if /XD, 11h if BU_LHRN not in FMS

**Format of DOS 4 x disk buffer hash chain head (array, one entry per chain):**

Offset	Size	Description
0	WORD	MS-DOS page number up with 0x1000 is resident, 1 if not in EMS
2	DWORD	sector of most recently used buffer header. All buffers on this chain are in the same segment
06h	BYTE	number of dirty buffers on this chain
07h	BYTE	reserved 00h
Note: $1 \leq N \leq 255$ reserved in 1 to chain N where N is the sector's address modulo number of disk buffer chain heads. $N \text{DIRTY} = 0 \leftrightarrow N \leq N \text{DIRTY} - 1$		
This structure is in main memory even if the buffers are in EMS		

**Format of DOS 4 0 6 0 disk buffer:**

Offset	Size	Description
00h	WORD	forward pointer (offset 00h) to next least recently used buffer
02h	WORD	backward pointer (offset 02h)
04h	BYTE	drive: 0=A 1=B, etc. FFh if not in use
05h	BYTE	<ul style="list-style-type: none"> <li>bit 7: remote buffer</li> <li>bit 6: buffer dirty</li> <li>bit 5: buffer has been referenced</li> <li>bit 4: search data buffer (only valid if remote buffer)</li> <li>bit 3: sector in data area</li> <li>bit 2: sector in a directory, either root or subdirectory</li> <li>bit 1: sector in FAT</li> <li>bit 0: reserved</li> </ul>
06h	DWORD	logical sector number
0Ah	BYTE	number of copies to write
		for FAT sectors, same as number of FATs
		for data and directory sectors, usually 1
0d	WORD	number of copies to write for FAT sectors
0f	DWORD	MS-DOS File Parameter Index (see 2f-32h)
1	WORD	number of copies to write for sectors above
2	BYTE	reserved
3		buffered data

Note: For DOS 4.x, all buffered sectors  $(x \neq 0)$  the sector's value computed as the sum of  $(x \times \text{DIRTY} + \text{DIRTY}) \times \text{DIRTY}$  on the same doubly linked circular chain; for DOS 5.x only a single circular chain exists.

This structure is in main memory even if the buffers are in the segment using the same forward buffers in the chain.

See BUFFERS.C in Chapter 8.

**Format of DOS 5+ disk buffer info:**

Offset	Size	Description
00h	WORD	MS-DOS page number used in buffer header, may be in HMA (see above)
04h	WORD	offset of DOS 5.x sector-based sectors, see offset 00h (points directly at the only buffer chain)
06h	DWORD	pointer to sector-based sectors, several not present
0Ah	WORD	number of logical sectors (may be zero, the value in BUFFERS.SYS)

0Ch	BYTE	buffer location 00h base memory, no workspace buffer 01h HMA, workspace buffer in base memory
01h	DWORD	pointer to core segment workspace buffer in base memory
11h	3 BYTES	apparently unused
14h	WORD	unknown
16h	BYTE	unknown counter, <i>is data</i>
17h	BYTE	temporary storage for use in memory allocation strategy during EMT
18h	BYTE	unknown counter
19h	BYTE	bit flags bit 0 unknown bit 1 SWITCHES - W specified in CONFIG.SYS don't auto-load WINA20.386 when MS Windows 3.0 starts bit 2 in EXEC state = 21 4B05h unknown offset - read only during 21 4B05h
1Ah	WORD	bit 0 set if CMB/MC Bch has asked to normal MC Bch
1Ch	BYTE	minimum paragraphs of memory required by program being EXEC'ed
1Dh	WORD	segment of Fast Memory (physical memory blocks of 1E11h or 100K memory chain in base 640k only) the first CMB/MC Bch is usually at 0E11h looking out into memory with a DOS-owned memory block
21h	WORD	paragraph of start of most recent MC Bch chain search

**Format of IFS driver list:**

Offset	Size	Description
00h	DWORD	pointer to next driver header
04h	8 BYTES	IFS driver name, blank padded (as used by THEMS command)
0Ch	4 BYTES	unknown
10h	DWORD	pointer to IFS utility function's entry point (see below) call with ES:BX pointer to IFS request (see below)
14h	WORD	offset to header's segment of driver entry point <i>possibly more</i>

Call IFS utility function entry point with

AH	20h	miscellaneous functions
AI	= 00h	get date
Returns:		
	CX	year
	DH	month
	DL	day
AI	= 01h	get process ID and computer ID
Returns:		
	BX	current PSP segment
	DX	active network machine number
AI	= 05h	get file system info
Returns:	ES:DI	pointer to 16-byte info buffer
		buffer filled
<b>Offset</b>	<b>Size</b>	<b>Description</b>
00h	2 BYTES	unused

8	WORD	number of SEFs actually counts only the first two file table arrays
04	WORD	number of FCB table entries
06	WORD	number of protected FCBs
8	6 BYTES	unused
	WORD	largest sector size supported
	AI - 00h get	machine name
		FS:DI                    pointer to 18 byte buffer for name
	<b>Returns</b>	buffer filled with name starting at offset 02h
	AI - 08h get	sharing retry count
	<b>Returns</b>	BX                    sharing retry count
	M   other	
	<b>Returns</b>	CI                    CI set
	MI 20h	
	AI 00h	
	M   other	
	<b>Returns</b>	CI                    CI set
	MI 20h get	redirection state
	BI                    CI set	
	<b>Returns</b>	BI                    state
	MI 22h	appears to be a time calculation
	AI 00h	
		CI set
	MI 23h	
	MI 24h	
	FS:SI	
	FS:DI	
	<b>Returns</b>	ZI                    ZI
	MI 25h	
	FS:SI	
	FS:DI	
	<b>Returns</b>	BI                    BI
	MI 26h	
	<b>Returns</b>	
	MI 27h	
	MI 28h	

**Note** IFS drivers which do not wish to respond to requests 70h or 74h-78h may pass them out to the default handler pointed at by 11:1-37h.

#### Format of IFS request block:

Offset	Size	Description
00h	WORD	total size in bytes of request
02h	BYTE	class of request



		02b unknown
		03b redirection
		04b unknown
		05h file access
		06h convert error code to string
		07b unknown
03h	WORD	returned IOCTL error code
05h	BYTE	IFS driver exit status
		00h success
		01b unknown
		02b unknown
		03b unknown
		04b unknown
		FFh internal failure
		unknown
06b	16 BYTES	
<b>—request class 02h—</b>		
16h	BYTE	function code
		04b unknown
17b	BYTE	apparently unused
18b	DWORD	pointer to unknown item
1Cb	DWORD	pointer to unknown item
20b	2 BYTES	unknown
<b>—request class 03h—</b>		
16h	BYTE	function code
17b	BYTE	unknown
18b	DWORD	pointer to unknown item
1Cb	DWORD	pointer to unknown item
22b	WORD	unknown returned value
24b	WORD	unknown returned value
26b	WORD	unknown returned value
28b	BYTE	unknown returned value
29b	BYTE	apparently unused
<b>—request class 04h—</b>		
16b	DWORD	unknown pointer
1Ab	DWORD	unknown pointer
<b>—request class 05h—</b>		
16h	BYTE	function code
		01h flush disk buffers
		02h get disk space
		03b MKDIR
		04b RMDIR
		05b CHDIR
		06b delete file
		07b rename file
		08h search directory
		09h file open/create
		0Ah LSEEK
		0Bh read from file
		0Ch write to file
		0Dh lock region of file

01h commit close file  
 01h get set file attributes  
 10h printer control  
 11h unknown  
 12h process termination  
 13h unknown

#### —class 05h function 01h—

17h 7 BYT  
 18h 1 DWORD  
 20h 4 BYT  
 26h BYT  
 27h BYT

unknown  
 unknown pointer  
 unknown  
 unknown  
 unknown

#### —class 05h function 02h—

17h 7 BYT  
 18h 1 DWORD  
 20h 4 BYT  
 26h WORD  
 28h WORD  
 2Ah WORD  
 2Ch WORD  
 2Eh BYT  
 2Fh BYT

unknown  
 unknown pointer  
 unknown  
 returned total clusters  
 unarmored sectors per cluster  
 returned bytes per sector  
 unarmored available clusters  
 unknown returned value  
 unknown

#### —class 05h functions 03h,04h,05h—

17h 7 BYT  
 1 1 DWORD  
 4 4  
 6 1 DWORD

unknown  
 unknown  
 unknown  
 pointer to destination region

#### —class 05h function 06h—

17h 7 BYT  
 18h 1 DWORD  
 20h 4 BYT  
 26h WORD  
 28h DWORD

unknown  
 unknown pointer  
 unknown  
 attribute value  
 pointer to filename

#### —class 05h function 07h—

17h 7 BYT  
 18h 1 DWORD  
 20h 4 BYT  
 26h WORD  
 28h DWORD  
 2Ch 1 DWORD  
 2Eh 1 DWORD

unknown  
 unknown  
 unknown  
 unknown  
 unknown  
 pointer to destination file spec

#### —class 05h function 08h—

18h 1 DWORD  
 20h 4 BYT  
 26h BYT  
 28h DWORD  
 2Ch WORD  
 2Eh DWORD

unknown  
 unknown  
 00h FINDFIRST  
 01h FINDNEXT  
 pointer to FindFirst search data + 01h if FINDNEXT  
 search attribute of FINDFIRST  
 pointer to file spec of FINDFIRST

#### —class 05h function 09h—

17h 7 BYT

unknown

1Eh	DWORD	<i>unknown pointer</i>
22h	DWORD	pointer to HFS open file structure (see below)
26h	WORD	<i>unknown</i>
28h	WORD	<i>unknown</i>
2Ah	4 BYTES	<i>unknown</i>
2Eh	DWORD	pointer to filename
32h	4 BYTES	<i>unknown</i>
36h	WORD	file attributes on call
		<i>unknown returned value</i>
38h	WORD	<i>unknown returned value</i>

—class 05h function 0Ah—

17h	7 BYTES	<i>unknown</i>
1Eh	DWORD	<i>unknown pointer</i>
22h	DWORD	pointer to HFS open file structure (see below)
26h	BYTE	sock type: 02h: from end
28h	DWORD	offset on call
		returned new absolute position

—class 05h functions 0Bh, 0Ch—

17h	7 BYTES	<i>unknown</i>
1Eh	DWORD	<i>unknown pointer</i>
22h	DWORD	pointer to HFS open file structure (see below)
28h	WORD	number of bytes to transfer
		returned bytes actually transferred
2Ah	DWORD	transfer address

—class 05h function 0Dh—

17h	7 BYTES	<i>unknown</i>
1Eh	DWORD	<i>unknown pointer</i>
22h	DWORD	pointer to HFS open file structure (see below)
26h	BYTE	<i>not used</i>
27h	BYTE	<i>apparently unused</i>
28h	WORD	<i>unknown</i>
2Ah	WORD	<i>unknown</i>
2Ch	WORD	<i>unknown</i>
2Eh	WORD	<i>unknown</i>

—class 05h function 0Eh—

17h	7 BYTES	<i>unknown</i>
1Eh	DWORD	<i>unknown pointer</i>
22h	DWORD	pointer to HFS open file structure (see below)
26h	BYTE	00h: common file
		01h: close file
27h	BYTE	<i>apparently unused</i>

—class 05h function 0Fh—

17h	7 BYTES	<i>unknown</i>
1Eh	DWORD	<i>unknown pointer</i>
22h	4 BYTES	<i>unknown</i>
26h	BYTE	02h: GET attributes
		03h: PUT attributes
27h	BYTE	<i>apparently unused</i>
28h	17 BYTES	<i>unknown</i>
34h	WORD	<i>variable attribute</i>
36h	DWORD	pointer to filename

136	WORD	GEI) unknown returned value
	WORD	GEI) unknown returned value
	WORD	GEI) unknown returned value
405	WORD	GEI) unknown returned value
420	WORD	PI F) new attributes
		GEI) returned attributes

**class 05h function 10h—**

	7 BYTES	unknown
	DWORD	unknown pointer
	DWORD	1 + 0x110HN + 0x01k + 0x11 + 0x110000
	WORD	unknown
	DWORD	unknown pointer
	WORD	unknown
	BYTE	unknown
3	BYTES	substitution
		01h get printer setup
		1 = unknown
		2 = unknown
		3 = unknown
		def. unknown
		07h unknown
		21h set printer setup

**class 05h function 11h—**

	BYTES	unknown
11h	DWORD	unknown pointer
7	DWORD	pointer to IF5 open file structure—see below
8	BYTES	substitution
	BYTE	apparently unused
	BYTES	+
	WORD	unknown
	WORD	unknown
	WORD	unknown
	WORD	unknown
	WORD	unknown

**class 05h function 12h—**

17h	15 BYTES	apparently unused
8	WORD	PN) + + +
20h	BYTE	Type of process termination
20h	BYTE	apparently unused

**class 05h function 13h—**

17h	15 BYTES	apparently unused
20h	WORD	PN) segment

**—request class 06h—**

00	DWORD	unknown 15 constants, some corresponding to other code at 03h
13h	BYTE	unknown returned value
13h	BYTE	unused

**—request class 07h—**

00	DWORD	unknown 15 constants, some corresponding to other code at 03h
13h	BYTE	apparently unused
00	BYTES	unknown 15 constants

**Format of IF5 open file structure.**

Offset	Size	Description
00h	WORD	unknown
02h	WORD	device into word
04h	WORD	file open mode

00h	WORD	<i>unknown</i>
08h	WORD	file attributes
0Ah	WORD	owner's network machine number
0Ch	WORD	owner's PSP segment
0Fh	DWORD	file size
12h	DWORD	current offset in file
16h	WORD	file time
18h	WORD	file date
1Ah	11 BYTES	filename in ECB format
16h	WORD	<i>unknown</i>
27h	WORD	real value of SE address on word of line (a file's segment & 1000h)
29h	3 WORDS	network info from SE 1
2Fh	WORD	<i>unknown</i>

#### Format of one item in DOS 4+ list of special program names:

Offset	Size	Description
00h	BYTE	length of string (00 is end of list)
01h	N BYTES	name in format name.ext
N	2 BYTES	DOS version required for program (major in low byte, 21-30h; 21-1221h)

#### —DOS 4 only—

N-2	BYTE	index of item to return (file size or 0000-FFFF hex)
-----	------	--

**Note** If the name of a file is executable for the program making the DOS "get version" call that is one of its names, this list DOS returns the specified version value that the true version might be.

### INT 21h Function 53h

**DOS 2+**

#### Translate Bios Parameter Block To Drive Parameter Block

Converts the information of a Drive Parameter Block from the information of the given BIOS Parameter Block.

#### Call With

AH	53h
DS:SI	pointer to BIOS Parameter Block (see below)
ES:BP	pointer to buffer of Drive Parameter Block (see 71-72 for format)

#### Returns

ES:BP buffer filled

**Note** For DOS 3, if a cluster at which the start sector is set to 0000h, and the number of clusters is set to FFFh, unknown.

#### Format of BIOS Parameter Block:

Offset	Size	Description
00h	WORD	number of bytes per sector
02h	BYTE	number of sectors per cluster
03h	WORD	number of reserved sectors at start of disk
05h	BYTE	number of FATs
06h	WORD	number of parameters of directory
08h	WORD	total number of sectors for DOS 4+, set to zero if partition > 32M (the set DWORD at 15h to actual number of sectors)
0Ah	BYTE	media ID byte
0Bh	WORD	number of sectors per FAT

**DOS 3+---**

0Dh	WORD	number of sectors per track
0Eh	WORD	number of heads
0Fh	WORD	number of hidden sectors
10h	WORD	reserved

**---DOS 4+---**

5h	WORD	total number of sectors; 1 word at 08h contains zero
6h	BYTE	unknown
7h	WORD	number of cylinders
21h	BYTE	device type
22h	WORD	device attributes: removable or not, etc.

**INT 21h Function 55h****DOS 2+****CREATE CHILD PSP**

Creates a child PSP in the Segment Prefix of the specified amount of available memory, and place it at a given location.

**Call With**

AH	55h
DX	segment at which to create new PSP
SI	0005h to 0007h: place in memory size (best at 0002h)

**Returns**

CF = 0 on success

**Notes:** 1. This function creates a full PSP rather than a copy of the current PSP, the current PSP's segment prefix is set to the current PSP's segment prefix, no count for each inherited file is maintained.

DOS 2: The current PSP is set to the segment in DX.

DOS 3: "Segmented" file handles are marked as closed in the child PSP.

See Also: 21, 26h, 21\_50h

**INT 21h Function 56h****DOS 2+****"RENAME" RENAME FILE**

Attempts to rename a file, may the file be source, destination or a directory, wildcards in both source and destination when invoked via 21\_51000h.

**Call With**

AH	56h
DS:DX	source: ASM IZ location of existing file (no wildcards, but see below)
ES:SI	pointer to ASM IZ new filename (no wildcards)
DI	attribute mask: server call only: see below

**Returns**

CF = 0 on success

CF = error code

AX = error code: 02h, 03h, 05h, 11h: see 71\_590

**Notes:** This function allows a file to be moved between different locations on a single logical volume.

Open files should not be renamed.

DOS 3: allows renaming of directories.

For DOS 3: wildcards are allowed: invoked via 21\_51000h, in which case error 12h (no more files) is returned on success, and both source and destination specs must be canonical (as returned by 21\_60h). Wildcards in the dest. are not allowed with the corresponding char of each source file being renamed. Under DOS 3X, the old wildcard in the destination wildcard is "\*" or equivalent. When invoked via 21\_51000h, only those files matching the attribute mask in DI are renamed.

See Also: 21/17b, 21/5D00h, 21/60h

## INT 21h Functions 5702h-5705h

DOS 4.x, OS/2, Chicago

### EXTENDED FILE ATTRIBUTE FUNCTIONS

In the DOS 4.x, OS/2, and Chicago 21/5702 sets the Extended Attributes (EAs) for a file, corresponding to the OS/2 `DosQueryPathInfo` and `DosSetPathInfo` API functions. Likewise, 21/5703 sets the EAs corresponding to `DosSetFileTime` and `DosSetFileInfo` Programs, but from the normal MS-DOS 21/57 Get/Set File Date/Time functions, which do not perform any OS/2 EA functions. For more information, see chapter 4 and also the register list below.

In DOS 4.x only, i.e., not DOS 5 or 6, 21/5702 and 21/5703 also set the file EA 21/5704, which has been reported in DOS 4.x as both EA function `OS/2 File EA Zero Length` and EA Set/Get Extended Attributes 21/5704, 21/5705 and 21/5704 as a call to `FILE_ATTRIBUTE_HIDDEN`. Under "Chicago" (OS/2 Windows) 21/5704 and 21/5705 will be the `FILE_ATTRIBUTE_READONLY` the Win32 kernel function and `SetFileTime` API function. Preliminary Chicago documents also marks 21/5705 bit function 21/13 as "reserved" with skip pointer over the two OS/2 EA functions.

## INT 21h Function 5D00h

DOS 3.1+

### SERVER FUNCTION CALL

Execute a specified INT 21h call using the sharing rules for the specified network machine number and process ID.

#### Call With

AX = 5D00h  
DS:DX = pointer to DOS parameter list (see below)  
DIP1 contains all register values for a call to INT 21h

#### Returns

as appropriate for function to be invoked

**Notes:** This function does not check the specified name for API on, off, or a network with the system.

The DOS call executes using specified computer ID and process ID, sharing delay (oops are skipped) a speed scoring mode (enabled with only for `DETECT` (21/41h) and `RENAME` (21/56h)), and an extra file attribute pointer is enabled for `OPEN` (21/33h), `DELETE` (21/41h), and `RENAME` (21/56h).

Functions which use the same register set as in other cases, as returned by 21/60h when invoked by this method, this is apparently to prevent multi-lop file forwarding.

See Also: 21/33h, 21/41h, 21/56h, 21/60h

### Format of DOS parameter list.

Offset	Size	Description
00h	WORD	AX
02h	WORD	BX
04h	WORD	CX
06h	WORD	DX
08h	WORD	SI
0Ah	WORD	DI
0Ch	WORD	DS
0Eh	WORD	ES
10h	WORD	reserved 0
12h	WORD	computer ID (0000h for current system)
14h	WORD	process ID (PSP segment on specified computer)

**Note:** Under Windows Extended mode, the computer ID is the virtual machine ID, see 21/68h.

**INT 21h Function 5D01h****DOS 3.1+****COMMIT ALL FILES FOR SPECIFIED COMPUTER/PROCESS**

This function commits and updates the directory entry for each file which has been written to since opening or the last commit.

**Call With**

AX 5D01h  
 DS:DX pointer to DOS parameter block (see 21 5D00h) only computer ID and process ID fields used

**Returns**

CF set on error  
 AX error code (see 21 50h)  
 CF clear if successful

**Notes:** This function commits and updates directory entries for each file which has been written to if the file is remote. It calls 21 110'h.

**See Also:** 21 50h, 21 51h, process ID is stored but ignored under DOS 3.3

**See Also:** 21 00h, 21 08h, 21 110'h

**INT 21h Function 5D02h****DOS 3.1+****SHARE.EXE CLOSE FILE BY NAME**

Close a file given its fully qualified name.

**Call With**

AX 5D02h  
 DS:DX pointer to DOS parameter block (see 21 5D00h) only DS:DX, computer ID, and process ID used  
 DI/ES:DI/DS pointer to ASCII name of file to close

**Returns**

CF set on error  
 AX error code (see 21 50h)  
 CF clear if successful

**Notes:** This function returns an error if SHARE is loaded. It calls SysfileTable 28h (see 21 52h).

The name must be a maximum of 255 bytes; such as is returned by 21 60h.

**See Also:** 21 40h, 21 5D03h, 21 5D04h, 21 60h

**INT 21h Function 5D03h****DOS 3.1+****SHARE.EXE CLOSE ALL FILES FOR GIVEN COMPUTER**

Close all files which were opened using a particular network machine number.

**Call With**

AX 5D03h  
 DS:DX pointer to DOS parameter block (see 21 5D00h) only computer ID used

**Returns**

CF set on error  
 AX error code (see 21 50h)  
 CF clear if successful

**Note:** This function returns an error if SHARE is loaded. It calls SysfileTable 30h (see 21 52h).

**See Also:** 21 5D02h, 21 5D04h



**INT 21h Function 5D04h****DOS 3.1+****SHARE.EXE - CLOSE ALL FILES FOR GIVEN PROCESS**

Close all files which were opened by a particular process.

**Call With**

AX            5306h  
 DS:DI       pointer to DOS parameter structure (SI=0000) (SI and DI  
 ID and process ID fields used)

**Returns**

CF set on error  
 AX        error code (see 21/506h)  
 CF clear if successful

**Note** This function is only available on systems that have SHARE installed. See the section on 21/52h.

**See Also:** 21/5102h, 21/5103h

**INT 21h Function 5D05h****DOS 3.1+****SHARE.EXE - GET OPEN FILE LIST ENTRY**

Return the contents of a file entry of a share record. A share record is a data structure.

**Call With**

AX            5105h  
 DS:DI       pointer to DOS parameter structure (SI=0000)  
 DDI & BX     index of sharing record (see 21/52h)  
 DDI & CX     index of SEI in sharing record's SEI list

**Returns**

CF clear if successful  
 FS:DI       pointer to AMIZ filename  
 BX        network machine number of SEI's owner  
 CX        number of locks held by SEI's owner  
 CF set if either index out of range  
 AX        0012h (no more files)

**Notes** This function is only available on systems that have SHARE installed. See the section on 21/52h.

**File** This function is implemented in the SHARE.SYS file. See the section on 21/60h.

**See Also:** 21/5102h, 21/60h

**INT 21h Function 5D06h****DOS 3.0+****GET ADDRESS OF DOS SWAPPABLE DATA AREA**

Return address of the DOS swappable data area. This function is restored to the BIOS by reentering.

**Call With**

AX            5306h

**Returns**

CF set on error  
 AX        error code (see 21/506h)  
 CF clear if successful

**File** This function is implemented in the DOS.SYS file. See the section on 21/60h.

**File** This function is implemented in the DOS.SYS file. See the section on 21/60h.

**File** This function is implemented in the DOS.SYS file. See the section on 21/60h.

Notes: 1. `INT 21h` and `INT 24h` are used in conjunction with the `INT05` flag (see 21-34h) to determine when it is safe to enter DOS from a TSR.

2. `INT 21h` and `INT 24h` may be the use of functions 50h-81h from `INT 28h` under DOS 2.1 by forcing use of correct stack.

3. `INT 21h` and `INT 24h` are used to enter DOS unless DOS is in a critical section delimited by 2A-80h and 2A-81h-82h.

4. `INT 21h` and `INT 24h` should be used instead of this function: the DOS 5+ swappable data area is also described under that call.

5. `INT 21h` and `INT 24h` may determine the SDA for their use by examining the byte at offset 04h of `INT 21h` or `INT 24h` (see `INT 21h`, AX 1205h). A value of 00h indicates a DOS 3.x SDA, a value of 80h+04h indicates a 3.x SDA in error. See `PSPIFN.C` in Chapter 4. See Also: 21-5100h; 2A-80h; 2A-81h; 2A-82h.

### Format of DOS 3.10 3.30 Swappable Data Area:

Offset	Size	Description
11	4 WORDS	critical error list of offsets which need to be pushed to enabled critical section calls (see 2A-80h); not actually part of the SDA
04	BYTE	critical error flag
05	BYTE	<code>INT05</code> flag (count of active <code>INT 21h</code> calls)
07	15 H	15 words which represent critical error occurred on 11h
08	BYTE	flags of last error
09	WORD	extended error code of last error
0a	BYTE	suggested action for last error
0c	BYTE	class of last error
0d	DWORD	ESI/EI pointer for last error
0e	DWORD	current DDA
0f	WORD	current PSP
2b	WORD	stores SP across an <code>INT 24h</code>
4b	WORD	error code from last process termination (cleared after reading with 21/40h, e.g. by <code>COMMAND.COM</code> )
10h	BYTE	current drive
17h	BYTE	extended break flag
<b>remainder need only be swapped if in DOS—</b>		
18h	WORD	value of AX on call to <code>INT 21h</code>
19	WORD	PSP segment for sharing network
1a	WORD	critical error number for sharing network (0000h=clear)
1c	WORD	critical error number block found when allocating memory
20	WORD	critical error number block found when allocating memory
22	WORD	critical error number block found when allocating memory
24	WORD	critical error number (perhaps used only during estimation)
26	WORD	unknown
28	BYTE	<code>INT 24h</code> returned flag
29	BYTE	bit flags for allowable actions on <code>INT 24h</code>
2a	BYTE	unknown flag
2b	BYTE	11h if Ctrl-Break termination, 00h otherwise
2c	BYTE	unknown flag
2d	BYTE	apparently not referenced
2f	BYTE	day of month
31	BYTE	month
30	WORD	year 1980
32	WORD	number of days since 1/1/1980

34h	BYTE	day of week (0 for Sunday)
35h	RTTI	working M/F pointer to M31-24 bit valid
36h	BYTE	safe to call INT 28h if nonzero
37h	BYTE	flag if nonzero; INT 74h abort turned into INT 24h if 1 (set only during process termination)
38h	26 BYTES	device driver request header (see 21_0002)
52h	DWORD	pointer to device driver entry point used in calling driver
56h	22 BYTES	device driver request header
6Ch	22 BYTES	device driver request header
82h	BYTE	type of PSP copy: 00h=simple for 21_26h 11h=make child
83h	RTTI	apparently not referenced by kernel
84h	3 BYTES	24 bit user number (see 21/30h)
87h	BYTE	CBM number (see 21_30h)
89h	2 BYTES	unknown
8Ah	6 BYTES	CIOR K5 transfer record (see 21_52h)
90h	BYTE	buffer for single byte LIO function
91h	BYTE	apparently not referenced by kernel
92h	128 BYTES	buffer for filename
112h	128 BYTES	buffer for filename
192h	21 BYTES	highest/lowest search data block (see 21_41h)
1A7h	32 BYTES	directory entry for found file
1C7h	81 BYTES	copy of current directory structure for drive being accessed
218h	11 BYTES	FCB format filename for device name comparison
226h	BYTE	apparently unused
224h	11 BYTES	wildcard distribution specification for filename (FCB format)
22Fh	2 BYTES	unknown
23Fh	WORD	unknown
246h	5 BYTES	unknown
248h	BYTE	extended FCB file attribute (highest search attribute)
239h	BYTE	type of PSP: 00h=regular 11h=extended
23Ah	BYTE	directory search attributes
23Bh	BYTE	file open mode
23Ch	BYTE	unknown flags, bits 0 and 4
23Dh	BYTE	unknown flag or counter
23Eh	BYTE	unknown flag
23Fh	BYTE	flag indicating how DOS interrupt was invoked (00h if loc, 1 INT 20h, INT 21h, 11h if service call INT 21 AX=5100h)
240h	BYTE	unknown
241h	BYTE	unknown flag
242h	BYTE	flag: 00h if read, 01h if write
243h	RTTI	unknown drive number
244h	RTTI	unknown
245h	RTTI	unknown flag or counter
246h	RTTI	file csn: ?1-0Ah: user mode flag: nonzero when user-owned direct function referred to existing block of file
247h	RTTI	unknown flag or counter
248h	RTTI	unknown flag or counter
249h	RTTI	type of process termination: 00h-03h (see 21_41h)
24Ah	RTTI	unknown flag
24Bh	RTTI	static sub which to replace first byte of deleted file's name (normally 15h but 00h as described under ?1_13h)
24Ch	DWORD	pointer to Drive Parameter Block for critical error invocation
24Dh	DWORD	pointer to stack frame containing user registers at INT 21h
24Eh	WORD	stores SP across INT 24h

	WORD	pointer to DOS Drive Parameter Block for unknown use
	WORD	unknown
	WORD	unknown temporary
	WORD	unknown flag, only low byte referenced
	WORD	unknown temporary
72	BYTE	Media ID byte returned by 21 1Bh, 21 1Ch
74	BYTE	apparently not referenced by kernel
76	DWORD	pointer to device header
78	DWORD	pointer to current SFT
79	DWORD	pointer to caller's directory structure for drive being accessed
7C	DWORD	pointer to caller's FCB
7E	WORD	offset in SFT to caller's FCB group identifier
7F	WORD	temporary storage for file handle
80	WORD	offset in FCB to caller's pointer to file, also see INT1 AH-26h
81	WORD	offset in DOS DS of first filename argument
82	WORD	offset in DOS DS of second filename argument
2800	WORD	offset of last component in pathname or FCBh
	WORD	offset of transfer address
84	WORD	relative cluster within file being accessed
86	WORD	absolute cluster number being accessed
88	WORD	current sector number
	WORD	current cluster number
	WORD	current offset in file DDI bytes per sector
8E	DWORD	unknown
	WORD	offset in file MDD bytes per sector
90	DWORD	current offset in file
	WORD	unknown
	WORD	unknown
	WORD	unknown
94	WORD	unknown
96	WORD	number of bytes appended to file
98	DWORD	pointer to unknown disk buffer
9A	DWORD	pointer to working SFT
9C	WORD	used by INT 21h dispatcher to store caller's BA
9E	WORD	used by INT 21h dispatcher to store caller's DS
	WORD	pointer to caller's system structure, also see INT 21h
9F	DWORD	pointer to prev call frame, offset 250h; if INT 21h terminated also switched to for duration of INT 24h
AA	21 BYTES	buffer for caller's buffer of a program operation see 21 41h
2A00	82 BYTES	directory entry for file being renamed
2A04	831 BYTES	critical error stack
1000	25 BYTES	scratch SFT
4380	284 BYTES	buffer for caller's structure, INT 25h INT 26h
4384	284 BYTES	character I/O stack, functions 01h through 0Ch
<b>—DOS 3.2, 3.3 only—</b>		
738h	BYTE	device driver lookahead flag, see 21 66h
740h	BYTE	apparently a drive number
742h	BYTE	unknown flag
744h	BYTE	unknown



3.1 The MS-DOS 5.0 program's registry also incorrectly states that the parameter list goes in DS:SI rather than in DS:DI.

See Also: 21-596

## INT 21h Function 5D08h

**DOS 4.x only**

### GET DOS SWAPPABLE DATA AREAS

25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49, 50, 51, 52, 53, 54, 55, 56, 57, 58, 59, 60, 61, 62, 63, 64, 65, 66, 67, 68, 69, 70, 71, 72, 73, 74, 75, 76, 77, 78, 79, 80, 81, 82, 83, 84, 85, 86, 87, 88, 89, 90, 91, 92, 93, 94, 95, 96, 97, 98, 99, 100, 101, 102, 103, 104, 105, 106, 107, 108, 109, 110, 111, 112, 113, 114, 115, 116, 117, 118, 119, 120, 121, 122, 123, 124, 125, 126, 127, 128, 129, 130, 131, 132, 133, 134, 135, 136, 137, 138, 139, 140, 141, 142, 143, 144, 145, 146, 147, 148, 149, 150, 151, 152, 153, 154, 155, 156, 157, 158, 159, 160, 161, 162, 163, 164, 165, 166, 167, 168, 169, 170, 171, 172, 173, 174, 175, 176, 177, 178, 179, 180, 181, 182, 183, 184, 185, 186, 187, 188, 189, 190, 191, 192, 193, 194, 195, 196, 197, 198, 199, 200, 201, 202, 203, 204, 205, 206, 207, 208, 209, 210, 211, 212, 213, 214, 215, 216, 217, 218, 219, 220, 221, 222, 223, 224, 225, 226, 227, 228, 229, 230, 231, 232, 233, 234, 235, 236, 237, 238, 239, 240, 241, 242, 243, 244, 245, 246, 247, 248, 249, 250, 251, 252, 253, 254, 255, 256, 257, 258, 259, 260, 261, 262, 263, 264, 265, 266, 267, 268, 269, 270, 271, 272, 273, 274, 275, 276, 277, 278, 279, 280, 281, 282, 283, 284, 285, 286, 287, 288, 289, 290, 291, 292, 293, 294, 295, 296, 297, 298, 299, 300, 301, 302, 303, 304, 305, 306, 307, 308, 309, 310, 311, 312, 313, 314, 315, 316, 317, 318, 319, 320, 321, 322, 323, 324, 325, 326, 327, 328, 329, 330, 331, 332, 333, 334, 335, 336, 337, 338, 339, 340, 341, 342, 343, 344, 345, 346, 347, 348, 349, 350, 351, 352, 353, 354, 355, 356, 357, 358, 359, 360, 361, 362, 363, 364, 365, 366, 367, 368, 369, 370, 371, 372, 373, 374, 375, 376, 377, 378, 379, 380, 381, 382, 383, 384, 385, 386, 387, 388, 389, 390, 391, 392, 393, 394, 395, 396, 397, 398, 399, 400, 401, 402, 403, 404, 405, 406, 407, 408, 409, 410, 411, 412, 413, 414, 415, 416, 417, 418, 419, 420, 421, 422, 423, 424, 425, 426, 427, 428, 429, 430, 431, 432, 433, 434, 435, 436, 437, 438, 439, 440, 441, 442, 443, 444, 445, 446, 447, 448, 449, 450, 451, 452, 453, 454, 455, 456, 457, 458, 459, 460, 461, 462, 463, 464, 465, 466, 467, 468, 469, 470, 471, 472, 473, 474, 475, 476, 477, 478, 479, 480, 481, 482, 483, 484, 485, 486, 487, 488, 489, 490, 491, 492, 493, 494, 495, 496, 497, 498, 499, 500, 501, 502, 503, 504, 505, 506, 507, 508, 509, 510, 511, 512, 513, 514, 515, 516, 517, 518, 519, 520, 521, 522, 523, 524, 525, 526, 527, 528, 529, 530, 531, 532, 533, 534, 535, 536, 537, 538, 539, 540, 541, 542, 543, 544, 545, 546, 547, 548, 549, 550, 551, 552, 553, 554, 555, 556, 557, 558, 559, 560, 561, 562, 563, 564, 565, 566, 567, 568, 569, 570, 571, 572, 573, 574, 575, 576, 577, 578, 579, 580, 581, 582, 583, 584, 585, 586, 587, 588, 589, 590, 591, 592, 593, 594, 595, 596, 597, 598, 599, 600, 601, 602, 603, 604, 605, 606, 607, 608, 609, 610, 611, 612, 613, 614, 615, 616, 617, 618, 619, 620, 621, 622, 623, 624, 625, 626, 627, 628, 629, 630, 631, 632, 633, 634, 635, 636, 637, 638, 639, 640, 641, 642, 643, 644, 645, 646, 647, 648, 649, 650, 651, 652, 653, 654, 655, 656, 657, 658, 659, 660, 661, 662, 663, 664, 665, 666, 667, 668, 669, 670, 671, 672, 673, 674, 675, 676, 677, 678, 679, 680, 681, 682, 683, 684, 685, 686, 687, 688, 689, 690, 691, 692, 693, 694, 695, 696, 697, 698, 699, 700, 701, 702, 703, 704, 705, 706, 707, 708, 709, 710, 711, 712, 713, 714, 715, 716, 717, 718, 719, 720, 721, 722, 723, 724, 725, 726, 727, 728, 729, 730, 731, 732, 733, 734, 735, 736, 737, 738, 739, 740, 741, 742, 743, 744, 745, 746, 747, 748, 749, 750, 751, 752, 753, 754, 755, 756, 757, 758, 759, 760, 761, 762, 763, 764, 765, 766, 767, 768, 769, 770, 771, 772, 773, 774, 775, 776, 777, 778, 779, 780, 781, 782, 783, 784, 785, 786, 787, 788, 789, 790, 791, 792, 793, 794, 795, 796, 797, 798, 799, 800, 801, 802, 803, 804, 805, 806, 807, 808, 809, 810, 811, 812, 813, 814, 815, 816, 817, 818, 819, 820, 821, 822, 823, 824, 825, 826, 827, 828, 829, 830, 831, 832, 833, 834, 835, 836, 837, 838, 839, 840, 841, 842, 843, 844, 845, 846, 847, 848, 849, 850, 851, 852, 853, 854, 855, 856, 857, 858, 859, 860, 861, 862, 863, 864, 865, 866, 867, 868, 869, 870, 871, 872, 873, 874, 875, 876, 877, 878, 879, 880, 881, 882, 883, 884, 885, 886, 887, 888, 889, 890, 891, 892, 893, 894, 895, 896, 897, 898, 899, 900, 901, 902, 903, 904, 905, 906, 907, 908, 909, 910, 911, 912, 913, 914, 915, 916, 917, 918, 919, 920, 921, 922, 923, 924, 925, 926, 927, 928, 929, 930, 931, 932, 933, 934, 935, 936, 937, 938, 939, 940, 941, 942, 943, 944, 945, 946, 947, 948, 949, 950, 951, 952, 953, 954, 955, 956, 957, 958, 959, 960, 961, 962, 963, 964, 965, 966, 967, 968, 969, 970, 971, 972, 973, 974, 975, 976, 977, 978, 979, 980, 981, 982, 983, 984, 985, 986, 987, 988, 989, 990, 991, 992, 993, 994, 995, 996, 997, 998, 999, 1000.

Call With

AX = 5D08h

Returns

CF = section error

AX = error code (see 21-996)

CF = clear if successful

DS:SI = pointer to swappable data area list (see below)

**Note:** Critical sections are swappable data areas (DOS 4.0) to be reentered unless it is in a critical section occurred by calls to 2A-80h and 2A-81h/82h.

Each entry in the DOS 4.0 swappable data area list is the SDA form at use by examining the byte at offset 01h. If the bit 0 is set, it is a DOS 4.0 SDA. If the bit 1 is set, it is a DOS 3.x SDA. A value of 00h indicates a DOS 3.x SDA. 01h indicates DOS 4.0/6.0 and other values are an error.

See Also: 21-5D06h, 2A-80h, 2A-81h, 2A-82h

### Format of DOS 4.x swappable data area list:

Offset	Size	Description
00h	WORD	count of data areas
02h	N BYTES	"count" copies of data area record
Offset	Size	Description
00h	DWORD	address
04h	WORD	length and type
		bit 15 set if swap always, clear if swap in DOS
		bits 14:0 = length in bytes

### Format of DOS 4.0/6.0 swappable data area:

Offset	Size	Description
1	STRUCTURE	critical section list of offsets which need to be patched to enabled critical section calls (see 2A-80h)
		not necessary if the SDA at offsets 00:00h but this list is still present for DOS 3.x compatibility
00h	BYTE	critical error flag
01h	WORD	DOS 4.0/6.0 or active INI 21 calls
02h	WORD	DOS 4.0/6.0 or active INI 21 calls
03h	BYTE	hex of last error
04h	WORD	extended error code of last error
06h	BYTE	suggested action for last error
07h	BYTE	class of last error
08h	DWORD	EB:DI pointer for last error
0Ch	DWORD	current DTA
10h	WORD	current PNP
12h	WORD	stores SP across an INI 23h

14h	WORD	return code from last process termination (cleared after reading with 21/41h) e.g. by COMMAND.COM
16h	BYTE	current drive
17h	BYTE	extended break flag
18h	2 BYTES	unknown
<b>—remainder need only be swapped if in DOS—</b>		
1Ah	WORD	value of AX on call to INT 21h
1Ch	WORD	PSP segment for sharing network
1Fh	WORD	network machine number for sharing network (0000h caller)
20h	WORD	best usable memory block found when allocating memory
22h	WORD	best usable memory block found when allocating memory
24h	WORD	best usable memory block found when allocating memory
26h	WORD	memory size in paragraphs (used only during initialization)
28h	WORD	unknown
2Ah	BYTE	unknown
2Bh	BYTE	unknown
2Ch	BYTE	unknown
2Dh	BYTE	unknown
2Eh	BYTE	unknown
2Fh	BYTE	apparently not referenced by kernel
30h	BYTE	day of month
31h	BYTE	month
32h	WORD	year (1980)
34h	WORD	number of days since 1/1/1980
36h	BYTE	day of week (0 for Sunday)
37h	BYTE	unknown
38h	BYTE	unknown
39h	BYTE	unknown
3Ah	30 BYTES	device driver request header (see 21/0002h)
58h	DWORD	pointer to device driver entry point (used in calling device driver request header)
5Ch	22 BYTES	device driver request header
72h	30 BYTES	device driver request header
90h	6 BYTES	unknown
96h	6 BYTES	IBM K5 transfer record (see 21/52h)
9Ch	2 BYTES	unknown
9Fh	128 BYTES	buffer for filename
11Fh	128 BYTES	buffer for filename
19Fh	21 BYTES	first indirect search data block (see 21/41h)
1B3h	32 BYTES	directory entry for found file
1D3h	88 BYTES	copy of current directory structure for drive being accessed
22Bh	11 BYTES	FCB format filename for device name comparison
236h	BYTE	unknown
237h	11 BYTES	wildcard destination specification for rename (FCB format)
242h	2 BYTES	unknown
246h	WORD	unknown
246h	5 BYTES	unknown
24Bh	BYTE	extended FCB file attributes (and first search attributes)
24Ch	BYTE	type of FCB (00h regular, FFh extended)
24Dh	BYTE	directory search attributes
24Eh	BYTE	file open mode
24Fh	BYTE	unknown or flag out
250h	BYTE	unknown flag or counter
251h	BYTE	unknown flag





2C Ah	WORD	used by INT 21h dispatcher to store caller's IV
2C Ch	WORD	used by INT 21h dispatcher to store caller's DS
2C Eh	WORD	temporary storage, while saving/restoring caller's registers
2D00h	DWORD	pointer to previous call frame (offset 764h of INT 21h register), also switched to for duration of INT 24h
2D04h	WORD	open mode action for 21 0C 00h
2D06h	BYTE	unknown; set to 00h in INT 21h dispatcher; 02h when a read is performed, and 01h or 03h by 21 0C 00h
2D76h	WORD	apparently unused
2D96h	DWORD	stored ESI for 21 0C 00h
2E10h	WORD	extended file open action code; see 21 0C 00h
2E12h	WORD	extended file open attributes; see 21 0C 00h
2E14h	WORD	extended file open file mode; see 21 0C 00h
2E16h	DWORD	pointer to filename to open; see 21 0C 00h
2E76h	WORD	unknown
2E96h	WORD	unknown
2E96h	BYTE	unknown
2E Ch	WORD	stores DS during call to List of Lists + 37h
2E E6h	WORD	unknown
2E E6h	BYTE	unknown
2E E6h	WORD	unknown but flags
2E 76h	DWORD	pointer to user-supplied filename
2E 76h	DWORD	unknown pointer
2E B6h	WORD	stores SS during call to List of Lists + 37h
2E D6h	WORD	stores SP during call to List of Lists + 37h
2E F6h	BYTE	flag; 00h if stack switched by calling List of Lists + 37h
3000h	21 BYTES	first list switch data for some files of a return operation (see 21 41h)
3150h	32 BYTES	directory entry for file being renamed
3350h	321 BYTES	critical error stack
4000h	384 BYTES	link stack; points to any greater than 00h of INT 25h/INT 26h
6000h	384 BYTES	character I/O stack; for various 01h through 06h
7800h	BY 1	kernel driver workable flag; see 21 04h
7816h	BTIE	apparently a driver number
7826h	BTIE	unknown flag
7836h	BTIE	unknown
7846h	BTIE	unknown
7856h	WORD	unknown
7866h	WORD	unknown
7876h	WORD	unknown
7886h	BTIE	unknown

## INT 21h Function SE01h SET MACHINE NAME

**DOS 3.1+ network**

Specifies the system's network machine name and number.

Call With:

AX	5E01h
CH	00h (undefine name; make it invalid)
	else define name
CL	name number
DS:DX	pointer to 16-character blank-padded ASCII name

See Also: 7E/7F00h

---

**INT 21h Function 5E04h**  
**SET PRINTER MODE**
**DOS 3.1+ network**

Specifies whether printer redirection is enabled or disabled.

**Call With**

CS           SI           DI           SI04h  
 AX           redirection list index (see 2F 5F09)  
 SI04h       word set if binary, clear if text (rdb expanded to blanks)

**Returns**

CF set on error  
 AX       error code (see 2F 5F09)  
 CF clear if successful

**Note** This function only works if the redirection list is set up. The redirection list is set up by the `SETUP` program. It is located at 2F 1111h, with AX pushed on stack.  
 See Also 2F 5F05h 2F 1111h

---

**INT 21h Function 5E05h**  
**GET PRINTER MODE**
**DOS 3.1+ network**

Retrieves the current printer redirection mode.

**Call With**

CS           SI05h  
 AX           redirection list index (see 2F 5F02h)

**Returns**

CF set on error  
 AX       error code (see 2F 5F09)  
 CF clear if successful  
 DX       printer mode (see 2F 5F04h)

**Note** This function only works if the redirection list is set up on the stack.  
 See Also 2F 5F04h 2F 1111h

---

**INT 21h Function 5F00h**  
**GET REDIRECTION MODE**
**DOS 3.1+ network**

Determines whether disk or printer redirection is currently enabled.

**Call With**

CS           SI00h  
 AX           redirection type: 03h printer, 04h disk drive

**Returns**

CF clear if successful  
 AX       redirection state: 00h on, 01h on

**Note** This function only works if the redirection list is set up on the stack.  
 See Also 2F 5F01h 2F 1111h

---

**INT 21h Function 5F01h**  
**SET REDIRECTION MODE**
**DOS 3.1+ Network**

Specifies whether disk or printer redirection is currently enabled.

**Call With**

AX	SEI01
BL	redirection type: 03h printer, 04h disk drive
BH	redirection state: 00h off, 01h on

**Returns**

CF	set on error
AX	error code (see 2F 506)
CF	clear if successful

**Note:** When redirection is off, the local device, if any, rather than the remote device is used. This function merely calls 2F 111Fh with AX on top of the stack.

**See Also:** 2F 5F00h, 2F 111Fh

**INT 21h Function 5F05h****DOS 4.x + Microsoft Networks****GET REDIRECTION LIST EXTENDED ENTRY**

Return the source and target of a given redirection, as well as its status and type.

**Call With**

AX	5F05
BX	redirection list index
DS:SI	pointer to buffer for ASCIZ source device name
ES:DI	pointer to buffer for destination ASCIZ network path

**Returns**

CF	set on error
AX	error code (see 2F 506)
CF	clear if successful
BH	device status flag: bit 0 clear if valid
BL	device type: 03h if printer, 04h if drive
CH	stored parameter value (user data)
BP	NFIBIOS local session number
DS:SI	buffer filled
ES:DI	buffer filled

**Notes:** If local session is used always set the redirection session number. However, if a local session is used, NFIBIOS is not applicable. In the case of a local drive, to correctly receive from errors, DS:SI = 0 and ES:DI = 2F 111Fh. If AX is on top of the stack.

**See Also:** 2F 5F06h, 2F 111Fh

**INT 21h Function 5F06h****Network****GET REDIRECTION LIST**

This function specifies to be a map to vector called 2F 5F02h get redirection list and 2F 5F05h get redirection list extended entry.

**Call With**

AX	5F06h
	additional arguments: at now, unknown

**Returns**

unknown

**Note:** This function merely calls 2F 111Fh with AX on top of the stack.

**See Also:** 2F 5F05h, 2F 111Fh

### INT 21h Function 5F07h ENABLE DRIVE

DOS 5+

Enables a drive that was previously temporarily disabled by setting the "valid" bit in the drive's Current Directory Structure.

#### Call With

AX     5F07h  
DI     drive number (0-A)

#### Returns

CF clear if successful  
CF set on error  
AX     error code (0fh see 21 50h)

See Also: 21 52h 21 5400h

### INT 21h Function 5F08h DISABLE DRIVE

DOS 5+

Disables a drive by clearing the "valid" bit in the drive's Current Directory Structure. For an alternate approach, see DRVNFUNC in Chapter 8.

#### Call With

AX     5F08h  
DI     drive number (0-A)

#### Returns

CF clear if successful  
CF set on error  
AX     error code (0fh see 21 50h)

See Also: 21 52h 21 5400h

### INT 21h Functions 5F32h 5F55h NAMED PIPES FUNCTIONS

LAN Manager

Microsoft LAN Manager (see 21 00h) provides Named Pipes and other services. Some of these calls are also supported by the Novell DOS Named Pipe Extension (Brows, VINES, and the OS 2 DOS box).

5F32h	DosQNmPipeInfo
5F33h	DosQNmPipeState
5F34h	DosQNmPipeState
5F35h	DosPeekNmPipe
5F36h	DosPeekNmPipe
5F37h	DosPeekNmPipe
5F38h	DosPeekNmPipe
5F39h	DosPeekNmPipe
5F3Ah	DosPeekNmPipe
5F3Bh	DosPeekNmPipe
5F3Ch	DosPeekNmPipe
5F3Dh	DosPeekNmPipe
5F3Eh	DosPeekNmPipe
5F3Fh	DosPeekNmPipe
5F40h	DosPeekNmPipe
5F41h	DosPeekNmPipe
5F42h	DosPeekNmPipe
5F43h	DosPeekNmPipe
5F44h	DosPeekNmPipe
5F45h	DosPeekNmPipe
5F46h	DosPeekNmPipe
5F47h	DosPeekNmPipe
5F48h	DosPeekNmPipe
5F49h	DosPeekNmPipe
5F4Ah	DosPeekNmPipe
5F4Bh	DosPeekNmPipe
5F4Ch	DosPeekNmPipe
5F4Dh	DosPeekNmPipe
5F4Eh	DosPeekNmPipe
5F4Fh	DosPeekNmPipe
5F50h	DosPeekNmPipe
5F51h	DosPeekNmPipe
5F52h	DosPeekNmPipe
5F53h	DosPeekNmPipe
5F54h	DosPeekNmPipe
5F55h	DosPeekNmPipe

These 21 5Fh calls are issued for example, by NETAPI.DLL in Windows, and by Microsoft's SQL Server. In the OS 2 2.0 DOS box, DDI, ACD, and FLEEBIRD make constant calls to 21 5F34h (DosPeekNmPipeState).

For more information, see *Computer Systems and Make Shivers* (The Undocumented LAN Manager and Named Pipes APIs on DOS and Windows) by *The Floppy Journal*, April 1993.



**INT 21h Function 6300h****DOS 2.25****GET DOUBLE BYTE CHARACTER SET LEAD BYTE TABLE**

IN: *dx* = *cs:dx* (0x) which specifies the range of characters which are the first half of a double byte character

Call With

AX = 6300h

Returns

CF clear if successful

DSI = pointer to lead byte table (see below for format)

all other registers except CS/SP and SS/SP destroyed

ZF set on error

AX = error code (01h) see 2f 59h

**Notes:** the US version of MS-DOS 3.30 (and later) and MS-DOS 3.40 (and later) treats this as a special case. If *dx* = 0000h, it returns a list of characters from the MS-DOS 6.00 Latin-1 and Windows-1252 tables to "DBCS.MY.DOS.2.21".  
See Also: 2f 07h 2f 08h 2f 09h 2f 0a01h 6301h

**Format of lead byte table entry:**

Offset	Size	Description
00h	2 1b	low byte of each of each of double byte chars
02h	2 1b	high byte of each of each of double byte chars
N	2 BYTES	00h/00h end flag

**INT 21h Function 6300h****Far East DOS 3.2+****GET DOUBLE BYTE CHARACTER SET LEAD BYTE TABLE**

IN: *dx* = *cs:dx* (0x) which specifies the range of characters which are the first half of a double byte character

Call With

AX = 6300h

Returns

AI = status

00h successful

DSI = pointer to DBCS table (see below)

all other registers except CS/SP and SS/SP destroyed

ZF not supported

**Notes:** the US version of MS-DOS 3.30 treats this as an unimplemented function returning AI=000h and returning immediately. the US version of DOS 3.40+ accepts this function but returns an empty list.

See Also: 2f 6301h 2f 65

**Format of DBCS table**

Offset	Size	Description
00h	2 1b	low byte of each of each of double byte chars
02h	2 1b	high byte of each of each of double byte chars
N	2 BYTES	00h/00h end flag

**INT 21h Function 6301h** **DOS 2.25, Far East DOS 3.2+****SET KOREAN (HANGEUL) INPUT MODE**

Specifies that DOS input characters will be returned partially formed, but not fully formed characters. This function sets the "interim console flag."

Call With

AX 6301h  
 DL new mode

00h successful (interim console flag set)  
 01h return partially formed interim characters also

Returns

AI status  
 00h successful  
 01h invalid mode

Note: see the Microsoft Windows Far East SDK.

See Also: 21 00h, 21 00b, 21 00c, 21 6300h, 21 6301h

**INT 21h Function 6302h** **DOS 2.25, Far East DOS 3.2+****GET KOREAN (HANGEUL) INPUT MODE**

Returns the current DOS input mode. See also set input mode function.

Call With

AX 6302h

Returns

AI status  
 00h successful  
 DL current input mode  
 00h successful (characters clear interim flag)  
 01h successful (characters set interim flag)  
 02h not supported

See Also: 21 00h, 21 00b, 21 00c, 21 6300h, 21 6301h

**INT 21h Function 64h** **DOS 3.2+****SET DEVICE DRIVER LOOKAHEAD FLAG**

Specifies that DOS input characters will be returned before the next character is typed. This device driver for input.

Call With

AH 64h  
 AI flag

00h successful (lookahead flag set)  
 before 21 01h 08h 0Ah  
 nonzero don't call driver function 5

Returns

AI set flag

Note: This function is called by DOS 3.3+ PRINT.COM

See Also: 21 00h, 21 00b, 21 00c, 21 00d, 21 00e, 21 00f, 21 010h, 21 011h, 21 012h, 21 013h, 21 014h, 21 015h, 21 016h, 21 017h, 21 018h, 21 019h, 21 01Ah, 21 01Bh, 21 01Ch, 21 01Dh, 21 01Eh, 21 01Fh, 21 020h, 21 021h, 21 022h, 21 023h, 21 024h, 21 025h, 21 026h, 21 027h, 21 028h, 21 029h, 21 02Ah, 21 02Bh, 21 02Ch, 21 02Dh, 21 02Eh, 21 02Fh, 21 030h, 21 031h, 21 032h, 21 033h, 21 034h, 21 035h, 21 036h, 21 037h, 21 038h, 21 039h, 21 03Ah, 21 03Bh, 21 03Ch, 21 03Dh, 21 03Eh, 21 03Fh, 21 040h, 21 041h, 21 042h, 21 043h, 21 044h, 21 045h, 21 046h, 21 047h, 21 048h, 21 049h, 21 04Ah, 21 04Bh, 21 04Ch, 21 04Dh, 21 04Eh, 21 04Fh, 21 050h, 21 051h, 21 052h, 21 053h, 21 054h, 21 055h, 21 056h, 21 057h, 21 058h, 21 059h, 21 05Ah, 21 05Bh, 21 05Ch, 21 05Dh, 21 05Eh, 21 05Fh, 21 060h, 21 061h, 21 062h, 21 063h, 21 064h, 21 065h, 21 066h, 21 067h, 21 068h, 21 069h, 21 06Ah, 21 06Bh, 21 06Ch, 21 06Dh, 21 06Eh, 21 06Fh, 21 070h, 21 071h, 21 072h, 21 073h, 21 074h, 21 075h, 21 076h, 21 077h, 21 078h, 21 079h, 21 07Ah, 21 07Bh, 21 07Ch, 21 07Dh, 21 07Eh, 21 07Fh, 21 080h, 21 081h, 21 082h, 21 083h, 21 084h, 21 085h, 21 086h, 21 087h, 21 088h, 21 089h, 21 08Ah, 21 08Bh, 21 08Ch, 21 08Dh, 21 08Eh, 21 08Fh, 21 090h, 21 091h, 21 092h, 21 093h, 21 094h, 21 095h, 21 096h, 21 097h, 21 098h, 21 099h, 21 09Ah, 21 09Bh, 21 09Ch, 21 09Dh, 21 09Eh, 21 09Fh, 21 0A0h, 21 0A1h, 21 0A2h, 21 0A3h, 21 0A4h, 21 0A5h, 21 0A6h, 21 0A7h, 21 0A8h, 21 0A9h, 21 0AAh, 21 0ABh, 21 0ACh, 21 0ADh, 21 0AEh, 21 0AFh, 21 0B0h, 21 0B1h, 21 0B2h, 21 0B3h, 21 0B4h, 21 0B5h, 21 0B6h, 21 0B7h, 21 0B8h, 21 0B9h, 21 0BAh, 21 0BBh, 21 0BCh, 21 0BDh, 21 0BEh, 21 0BFh, 21 0C0h, 21 0C1h, 21 0C2h, 21 0C3h, 21 0C4h, 21 0C5h, 21 0C6h, 21 0C7h, 21 0C8h, 21 0C9h, 21 0CAh, 21 0CBh, 21 0CCh, 21 0CDh, 21 0CEh, 21 0CFh, 21 0D0h, 21 0D1h, 21 0D2h, 21 0D3h, 21 0D4h, 21 0D5h, 21 0D6h, 21 0D7h, 21 0D8h, 21 0D9h, 21 0DAh, 21 0DBh, 21 0DCh, 21 0DDh, 21 0DEh, 21 0DFh, 21 0E0h, 21 0E1h, 21 0E2h, 21 0E3h, 21 0E4h, 21 0E5h, 21 0E6h, 21 0E7h, 21 0E8h, 21 0E9h, 21 0EAh, 21 0EBh, 21 0ECh, 21 0EDh, 21 0EEh, 21 0EFh, 21 0F0h, 21 0F1h, 21 0F2h, 21 0F3h, 21 0F4h, 21 0F5h, 21 0F6h, 21 0F7h, 21 0F8h, 21 0F9h, 21 0FAh, 21 0FBh, 21 0FCh, 21 0FDh, 21 0FEh, 21 0Fh

See Also: 21 01h, 21 00b, 21 00c, 21 5100h





**Format of filename terminator table:**

Offset	Size	Description
00h	WORD	table size, not counting 0's word
02h	BYTE	unknown 01h for MS-DOS 3.30-6.00
03h	BYTE	low set printable character set (0-9, a-z)
04h	BYTE	high set printable character set (A-Z, space)
05h	BYTE	unknown 00h for MS-DOS 3.30-6.00
06h	BYTE	bits excluded via sector size (1 = characters not as
07h	BYTE	list exclude 1 character (0 = none) - see are illegal
08h	BYTE	unknown 02h for MS-DOS 3.30-6.00
09h	BYTE	number of illegal terminal characters
0Ah	STRING	characters of extension table (0 = none)

**Note** Partially documented for DOS 5.0, but undocumented for earlier versions.

**INT 21h Functions 6520h to 6522h****DOS 4+****COUNTRY DEPENDENT CHARACTER CAPITALIZATION**

Capitalize a character or string using the capitalization rules for the current country.

**Call With**

AH = 65h  
 AL = function

20h capitalize character

DL = character to capitalize

**Returns:**

DL = capitalized character

21h capitalize string

DS:DX = pointer to string to capitalize

CX = length of string

22h capitalize ASCII string

DS:DX = source ASCII string to capitalize

**Returns**

CF set on error

AX = error code (see 21-50h)

CF clear if successful

**Note** This call is not supported for DOS 5.0. See section 5.1.4.2 for details on DOS 4.x.

**INT 21h Function 6523h****DOS 4+****Determine If Character Represents Yes/no Response**

Compare the specified character against the Y/N and N/O responses for the current country.

**Call With**

AX = 6523h  
 DI = character  
 DI = second character of double-byte character, if applicable

**Returns**

CF set on error

CF clear if successful

AX = type

00h no

01h yes

02h neither yes nor no

## INT 21h Functions 65A0h to 65A2h Country dependent Filename Capitalization

DOS 4+

### Call With

AH 65h  
 AL function  
 AHb capitalize filename character

### Returns:

AL capitalized character  
 AXb capitalization count of filename string  
 DS:DI pointer to filename string  
 CX length of string  
 AHb capitalization of filename  
 DS:DI pointer to ASCII filename to capitalize

### Returns

CF set on error  
 AX error code (see M 50b)  
 CF clear if successful

**Note:** This function is available only on DOS 4.00 and later. It is not available on DOS 3.31 and earlier.

## INT 21h Function 67h SET HANDLE COUNT

DOS 3.3+

This function sets the maximum number of handles that a process can have. This function is available only on DOS 3.31 and later.

### Call With

AH 67h  
 AL size of new file handle table for process

### Returns

CF clear if successful  
 CF set on error  
 AX error code (see M 50b)

**Note:** This function is available only on DOS 3.31 and later. It is not available on DOS 3.30 and earlier.

For file handle tables, the maximum number of handles is 20. On DOS 3.31 and later, the maximum number of handles is 255. This function sets the maximum number of handles for a new process allocated and the system.

Only the first 20 handles are copied to child processes in DOS 3.31-6.0.

**Increment:** This function is available only on DOS 3.31 and later. It is not available on DOS 3.30 and earlier.

**BUGS:** This function is available only on DOS 3.31 and later. It is not available on DOS 3.30 and earlier.

See Also: 21\_26h

## INT 21h Function 69h GET/SET DISK SERIAL NUMBER

DOS 4.0+

Determine or specify a disk's serial number and volume label.

**Call With**

AH	69h
AL	subfunction 00h get serial number 01h set serial number
BI	drive 0=default 1-3, 2+1 etc.
DS:DI	pointer to disk info (see below)

**Returns**

CF	set on error
AX	error code (see 2F 59b)
CF	clear if successful
AX	destroyed
DI	-00h but called with appropriate values from extended BPB
DI	-01h extended BPB on disk set to values from buffer

**Notes**

- Also if you check for error the location of a sector is returned in AX
- From 00h set serial # of a disk's extended BLP on disk
- This call does not work on network drives (error 0001h)
- Also, the first two bytes of the buffer is exactly opposite to 2F00h and 3F00h of the extended BPB on the disk

This function is supported under Novell NetWare versions 2.0A through 3.11; the returned serial number is set to zero if a DFR would display the volume label in the NetWare volume label and the file system is set to "FAT16".

The volume label which is read or set by this function is the one stored in the extended BPB on disks for serial # 00h-30540h rather than the special root directory entry used by the DFR command (COMMAND.COM use 2F 1F 1F mod'54) volume label.

See Also: 21 4400h (X-0800h Get Media ID)

**Format of disk info:**

Offset	Size	Description
00h	WORD	disk 1 ser
02h	DWORD	disk serial number (binary)
06h	11 bytes	volume label (NO NAME if none present)
11h	8 bytes	FILESYSTEMTYPE string "FAT12" or "FAT16"

**INT 21h Function 6Ah****DOS 4+****COMMIT FILE**

This call is a little bit undocumented (function 6Ah is DOS 5.0 and 6.0 but is not known whether the two functions are identical in DOS 4.x)

**Call With**

AH	6Ah
BX	file handle

**Returns**

CF	clear if successful
AH	68h
CF	set on error
AX	error code

See Also: 2F 68h

**INT 21h Function 6Bh****DOS 4.x only****UNKNOWN FUNCTION**

The purpose of this function is not known, but it appears to be related to a label file system.

**Call With**

**AH** 6Bh  
**AI** subfunction  
**DI** *optional*  
     *DOS pointer to current Directory Structure*  
     CX drive 1-A  
**SI** *optional*  
     *DOS pointer to unknown item*  
     CX *file handle*  
**ES** *optional*  
     *DOS pointer to current Directory Structure*  
     DI *unknown*  
     CX drive 1-A

**Returns**

**CF** set on error  
**AX** error code (see 21 50h)  
**CF** clear if successful

**Note** This call is passed through INT 21 112h with AX on top of the stack.  
**See Also:** 21 112h

**INT 21h Function 6Bh****DOS 5+****NULL FUNCTION**

This function performs no action and returns no information.

**INT 21h Functions 6Dh-6Fh****MS-DOS 5+ in ROM, OS/2****DOS IN ROM FUNCTIONS**

Access to ROM is available in MS-DOS 5.0 and later (see 21 6Dh through 21 6Fh) for a set of functions that are available to MS-DOS 5.0 and later (see 21 6Dh through 21 6Fh).

**6Dh** Find First ROM File  
**6Eh** Find Next ROM Program  
**6F00h** Get ROM Scan Start Address  
**6F01h** Set ROM Scan Start Address  
**6F02h** Get Exclusion Region List  
**6F03h** Set Exclusion Region List

For more information, see the interrupt listing disk (and Chappell's *DOS Internals*). To detect DOS in ROM, see indocumented 21 330h.

**File:** C:\DOS\IOPLIB.C (MS-DOS 5.0) or C:\DOS\IOPLIB2.C (MS-DOS 5.0+)

**File:** C:\DOS\IOPLIB.C (MS-DOS 5.0) or C:\DOS\IOPLIB2.C (MS-DOS 5.0+)

**INT 21h Functions 71h-72h****Chicago****CHICAGO LONG FILENAME FUNCTIONS**

Under Chicago, DOS 5.0, Windows 3.11, and NT 3.51 (and, DOS programs with access to our file system). A DOS program will be able to use the following functions where the subfunction is a subfunction code. DOS 5.0 and later (see 21 71h through 21 74h) because the old DOS Get Current Directory Name function (21 71h) returns a maximum of 255 characters, while the old directory identifiers returned by DOS must be larger than the maximum number of characters in a normal path name. The Chicago 21 74h Get Volume Information function to get the length of the path name.

The following are to be the documented INT 21h long filename functions, with their Win32 API equivalents:

7139h	CreateDirectory
713Ah	RemoveDirectory
713Bh	SetCurrentDirectory
7141h	DeleteFile
7142h	FileDirectoryAttributes
7143h	FileFindFirstFileAttributes
7147h	CreateDirectory
7141h	FileFindFile
714Fh	FindNextFile
7156h	MoveFile
716Ch	CreateFile and OpenFile

If the 21h INT function has a carry set (AH=1000h) indicates the function was not supported for example, the volume of a sector support long filenames. In this case continue by calling the old 21h INT version 1 (example 32a "24" table) carry set and AH=1000h call 21h 4.

A 21h INT long filename version 7140h function does not appear to be provided.

Another way, 8086 uses 21h 2 and RAX. Which is previous versions of DOS support only find first file find next file requires find file because the Win32 API supports multiple simultaneous file finds find first 21h 7141 and find next file which you must pass to find next 21h 7141 and which you must find file case with find file.

For more information see chapter 8 and eventually Microsoft volume 6 or for Chicago.

## INT 21h Function 80h

## European MS-DOS 4.0

### AEXEC - EXECUTE PROGRAM IN BACKGROUND

Asynchronously execute a program, creating a new process for it. Although this function and the known functions on the next several pages are in fact obsolete, I included them here only because of my knowledge that various early European versions of the more widely popular MS-DOS 4.0s and OEM versions with their various capabilities which were sold on a large number of systems between the mainstream versions 3.2 and 5. Some features of Windows 1.03S 2 including the New Executable (NE) file format first appear in European MS-DOS 4.0 Sec. 21-87 for the European MS-DOS install check.

### Call With:

AH	80h
CX	mode
	0000h place child in zombie mode or exit to preserve exit code
	0001h discard child process and exit code on termination
DS:DI	pointer to ASCII null program name
ES:BX	pointer to parameter block as for 21-4B04h.

### Returns:

CF	clear if successful
AX	Command Subgroup ID (CMD)
CF	set on error
AX	error code (see 21-59h)

**Notes:** This function is called by the DEEXEC command.

There is a system-wide limit of 32 processes.

See INT 21h Function 81h for details on how to specify a given process, whether directly or indirectly.

See INT 21h Function 82h for details on how to specify a command. See INT 21h Function 83h for details on how to specify a foreground process. See INT 21h Function 84h for details on how to specify a background process.

See INT 21h Function 85h for details on how to specify a command and a process. See INT 21h Function 86h for details on how to specify a command and a foreground process. See INT 21h Function 87h for details on how to specify a command and a background process.

See INT 21h Function 88h for details on how to specify a command and a process and a foreground process. See INT 21h Function 89h for details on how to specify a command and a process and a background process.

See INT 21h Function 8Ah for details on how to specify a command and a process and a foreground process and a background process. See INT 21h Function 8Bh for details on how to specify a command and a process and a foreground process and a background process and a foreground process.

See INT 21h Function 8Ch for details on how to specify a command and a process and a foreground process and a background process and a foreground process and a background process. See INT 21h Function 8Dh for details on how to specify a command and a process and a foreground process and a background process and a foreground process and a background process and a foreground process.

### INT 21h Function 81h "FREEZE" STOP A PROCESS

European MS-DOS 4.0

Freeze a process. See INT 21h Function 80h for general comments on European MS-DOS 4.0.

#### Call With

AH: 81h

CX: Process ID of head of command subtree.

#### Returns

CF: clear if successful

CF: set on error

AX: error code (no such process)

Note: A process which is frozen will not receive any signals which may not be until after it unblocks from an I/O operation.

See Also: AH: 82h, AH: 89h, AX: 8100h.

### INT 21h Function 82h "RESUME" RESTART A PROCESS

European MS-DOS 4.0

Restart a process. See INT 21h Function 80h for general comments on European MS-DOS 4.0.

#### Call With

AH: 82h

BX: flag 00h resume command subtree 01h only specified process

CX: Process ID of head of command subtree.

#### Returns

CF: clear if successful

CF: set on error

AX: error code (no such process)

See Also: INT 21h 81h.

### INT 21h Function 83h "PARTITION" GET/SET FOREGROUND PARTITION SIZE

European MS-DOS 4.0

Get/set foreground partition size. See INT 21h Function 80h for general comments on European MS-DOS 4.0.

#### Call With

AH: 83h

AI function  
 01h get size  
 02h set new size  
 BX new size in paragraphs

**Returns**

CF clear if successful  
 BX current size (function 000h) or old size (function 01h)  
 CF set on error  
 AX error code (01h-07h/0Dh; see ME 59f)

**Notes:** If the partition system is set to 0H/0h, no partition management is done and all memory allocation is compatible with DOS 3.2.

The partition size can be changed regardless of what use is being made of the choice; after subsequent allocations will follow the partition rules. (Only one of processes may allocate on the ground plane; a file could process all at background memory for then, for program error only.)

See Also: INT 21 Function 81h, 21\_4Ah

### INT 21h Function 8400h European MS-DOS 4.0 **"CREATMEM" CREATE A SHARED MEMORY AREA**

Create a shared memory object that may be accessed by multiple processes. See INT 21h Function 800h for general information on European MS-DOS 4.0.

**Call With**

AX 8400h  
 BX size in bytes (0000h-65536h)  
 CX flags  
   bit 0: zero initialize segment  
 DS:DX pointer to ASCII name (must begin with "SHAREMEM")

**Returns**

CF clear if successful  
 AX segment address of shared memory global object  
 CF set on error  
 AX error code (06h/09h; see INT 11h, INT 85h and INT 1Bh Function 86h)

**Notes:** Shared memory objects are created as special files; thus the restriction on the name.

On successful creation, the reference count is set to 1.

See Also: 21\_8401h, 21\_8402h

### INT 21h Function 8401h European MS-DOS 4.0 **"GETMEM" OBTAIN ACCESS TO SHARED MEMORY AREA**

Get address of previously created shared memory object that may be accessed by multiple processes.

**Call With**

AX 8401h  
 CX flags  
   bit 0: writeable segment (ignored by MS-DOS 4.0)  
 DS:DX pointer to ASCII name (must be same as "SHAREMEM")

**Returns**

CF clear if successful  
 AX segment address of shared memory global object  
 CX size in bytes  
 CF set on error  
 AX error code (06h) and name.

**Note:** This call increments the reference count for the shared memory area.  
 See Also: AX:8400h, AX:8402h

---

### INT 21h Function 8402h European MS-DOS 4.0

#### RELEASEMEM™ FREE SHARED MEMORY AREA

Releases the shared memory will no longer be used by the caller.

**Call With**

- AX: 8402h
- DX: validly segment address of shared memory area

**Returns**

- CF: clear if successful
- CF: set on error
- AX: error code (no such thing)

**Note:** This function is not available in MS-DOS 3.0 and MS-DOS 3.11.

See Also: 21:8400h, 21:8401h

---

### INT 21h Function 86h European MS-DOS 4.0

#### SETFILETABLE™ INSTALL NEW FILE HANDLE TABLE

Replaces the current file handle table with a new table. The new table must be a contiguous block of bytes the caller can open simultaneously.

**Call With**

- AX: 86h
- AX: total number of file handles in new table

**Returns**

- CF: clear if successful
- CF: set on error
- AX: error code (00h, 08h) (see AX:590h)

**Note:** Any currently open files are copied to the new table.  
 This function is not available in MS-DOS 3.0 and MS-DOS 3.11.

Using current open files

See Also: 21:26h, 21:67h

---

### INT 21h Function 87h European MS-DOS 4.0

#### GETPID™ GET PROCESS IDENTIFIER

Obtains the process ID of the caller. This function is used as a standard check for European MS-DOS 4.0.

**Call With**

- AX: 87h

**Returns**

- AX: Process ID (PID)
- DX: parent process PID
- CX: Command Subgroup ID (CSID)

**Note:** This function is not available in MS-DOS 3.0 and MS-DOS 3.11.  
 This function is not available in MS-DOS 3.0 and MS-DOS 3.11.

Obtains the process ID of the caller. This function is used as a standard check for European MS-DOS 4.0. See also AX:400h and check whether AX is nonzero on return.

See Also: 21:30h, 21:62h, 21:80h



**INT 21h Function 89h****European MS-DOS 4.0****SLEEP**

Suspend the calling process for the specified duration.

**Call With**

- AX = 89h
- CX = time in milliseconds or 0000h to give up time slice

**Returns**

- CF clear if successful
- CX = 0000h
- CF set on error
- AX = error code (no corrupted system call)
- CX = sleep time remaining

**Note:** The sleep interval is rounded up to the next higher increment of the scheduler. Execution may be extended further if other processes are waiting.

This call may be interrupted if signals are pending.

See *Microsoft Windows 3.11 Programmer's Reference*, Sec. 21, NDF.

See *Microsoft Windows 3.11 Programmer's Reference*, Sec. 21, NDF.

See *Microsoft Windows 3.11 Programmer's Reference*, Sec. 21, NDF.

See *Microsoft Windows 3.11 Programmer's Reference*, Sec. 21, NDF.

See Also: 21, 81h, INT 15, AX=1000h, 21-1000h

**INT 21h Function 8Ah****European MS-DOS 4.0****"CWAIT" WAIT FOR CHILD TO TERMINATE**

Wait for a child to terminate. This function suspends the calling process until the child terminates. The return code is returned. See INT 21h Function 80h for general comments on European MS-DOS 4.0.

**Call With**

- AX = 8Ah
- BI = range: 00h command subtree, 01h any child
- BF = suspend flag
  - 00h suspend if children exist but none are dead
  - 01h return if no dead children
- CX = Process ID of head of command subtree

**Returns**

- CF clear if successful
- AX = termination type
  - 00h normal termination
  - 01h aborted by Control C
  - 02h aborted by Ctrl-C or Ctrl-Break
  - 03h terminate and stay resident
  - 04h aborted by signal
  - 05h aborted by program error
- AI = return code from child or aborting signal
- BX = PID of child, 0000h if no dead children
- CF set on error
  - AX = error code (no child, interrupted system call)

See Also: 21, 4B04h, 21-4D6, 21-806, 21-8D6

**INT 21h Function 8Ch****European MS-DOS 4.0****SET SIGNAL HANDLER**

Set the signal handler for a child. It is a member of exception handling.

**Call With**

AH 8Ch  
 AI signal number (see below)  
 BI action (see below)  
 DS:DI pointer to signal handler (see below)

**Returns**

CF clear if successful  
 AI previous action  
 EBX pointer to previous signal handler  
 CF set on error  
 AX error code (if a valid SigNumber or Action): see AII-540b

**Note** A signal is sent to the most recently installed handler.

See Also: 21-800b

**Values for signal number:**

01h	SIG_NUL	no signal or user-defined interrupt key
08h	SIG_TERM	program termination
09h	SIG_PIPE	broken pipe
0Jh	SIG_USER1	reserved for user definition
0Kh	SIG_USER2	reserved for user definition

**Values for signal action:**

00h	SIG_DFL	terminate process on receipt
01h	SIG_IGN	ignore signal
02h	SIG_CGT	signal is accepted
03h	SIG_ERR	sender gets error
04h	SIG_ACK	acknowledge received signal and clear it, but don't change current setting

**Signal handler is called with:**

AI signal number  
 AI signal arguments

**Returns**

RET, CF set: terminate process  
 RET, CF clear, ZF set: user-defined interrupted system call with error  
 RET, CF clear, ZF set: start a user-defined system call  
 RET, CF clear, CF set: user-defined system call

**Note** The user-defined system call may perform a `CALL` or `JMP` to GOTO by resetting the stack pointer and jumping to the location specified by `DISPATCH` the signal by calling this function with BI=04h.

**INT 21h Function 8Dh****European MS-DOS 4.0****SEND SIGNAL**

Invoke the exceptional condition handler for the specified process.

**Call With**

AH 8Dh  
 AI signal number (see AII-8C h)  
 BI signal argument  
 BI action  
   00h send to entire command subtree  
   01h send only to specified process  
 DI Process ID

**Returns**

CF clear if successful  
 CF set on error

AX error code (01h-06h) see AH-59h

**Note** Error 06h may be returned if one or more of the affected processes have an error handler for the signal.

See Also: 21-80h

### INT 21h Function 8E00h

European MS-DOS 4.0

#### "SETPRI" - GET/SET PROCESS PRIORITY

Specify or determine the execution priority of the specified process or the process and all of its children.

#### Call With

AX 8E00h  
 BH 00h  
 BL action  
     00h set priority for command subtree  
     01h set priority for specified process only  
 CX Process ID  
 DH 00h  
 DI change in priority (00h to get priority)

#### Returns

CF clear if successful  
 DE process priority  
 DH destroyed  
 CF set on error  
 AX error code (01h) no such process (see 21-50h)

See Also: 21-80h

### INT 21h Function 93h

European MS-DOS 4.0

#### "PIPE" - CREATE A NEW PIPE

Create a communications channel which may be used to interprocess data and command exchanges.

#### Call With

AH 93h  
 CX size in bytes

#### Returns

CF clear if successful  
 AX read handle  
 BX write handle  
 CF set on error  
 AX error code (08h) (see AH-59h)

See Also: 21-3C; 21-3B; 21-40h; 21-84h

### INT 21h Function 95h

European MS-DOS 4.0

#### HARD ERROR PROCESSING

Specify whether hard (critical) errors should be automatically fail the system call or invoke an INT 24h.

#### Call With

AH 95h  
 AL new state  
     00h enabled  
     01h disabled, automatically fail hard errors

**Returns**

AX - previous setting

See Also: INT 20h

**INT 21h Function 99h****European MS-DOS 4.0****PBLOCK - BLOCK A PROCESS**

Set a timer to block a process until another process sends a signal or a timeout occurs.

**Call With**

AX - 99h

DS:BX - pointer to memory location of which to block

CX - time in tenths of seconds

DH - nonzero if interruptible

**Returns**

CF - clear if awaked by event

AX - 0000h

CF - set if unusual waking

ZF - set if timeout, clear if interrupted by signal

AX - nonzero

See Also: 21 9Ah, 21 0802h

**INT 21h Function 9Ah****European MS-DOS 4.0****PRUN - UNBLOCK A PROCESS**

Restart all processes waiting for the specified "break" signal.

**Call With**

AX - 9Ah

DS:DX - pointer to memory location of break processes to unblock

**Returns**

AX - number of processes awakened

ZF - set if no processes awakened

See Also: 21 99h, 21 0802h

**INT 28h****DOS 2+****DOS IDLE INTERRUPT**

The DOS idle interrupt is a special interrupt that MS-DOS has set up to function as a timeout in DOS. When the timeout occurs, the interrupt signals all DOS to allow the user to interact with the system. The interrupt is set up to occur when the program is waiting for user input.

**Called with**

SS:SP - top of MS-DOS stack for I/O functions

**Returns**

all registers preserved

**Notes:** INT 28h is a special interrupt that MS-DOS has set up to function as a timeout in DOS. When the timeout occurs, the interrupt signals all DOS to allow the user to interact with the system. The interrupt is set up to occur when the program is waiting for user input. The interrupt is set up to occur when the program is waiting for user input.

Call the interrupt by setting the interrupt vector in the INT 28h table. The interrupt vector is a handle which refers to the DOS device.



## Returns

**INT 2Ah Function 0300h****Network****CHECK DIRECT I/O**

Call with  $AX = 0300h$ .

## Call With

$AX = 0300h$

## Returns

1) Check if direct physical addressing (IN 13h) IN 25h) permitted.  
 2) See EXECUTE IOCTL.

## See

Also:

INT 13h, INT 25h, INT 26h, INT 500h

**INT 2Ah Function 04h****Network****EXECUTE NETBIOS REQUEST**

## Call With

$AX = 04h$

$SI =$

Word indicating retry request on error 09h, 12h, and 21h  
 0E) network  
 pointer to Network Control Block

## Returns

0000) Success

01) error

Al) error code

Note: If system 0E) invokes either INT 5Bh or INT 5Ch as appropriate.

See Also: 7A) 0A) 0A) 0900h, INT 5Bh, INT 5Ch

**INT 2Ah Function 0500h****Network****GET NETWORK RESOURCE AVAILABILITY**

Determine the available amounts of several important network resources.

## Call With

$AX = 0500h$

## Returns

1) Number of network nodes available

2) Number of networks

3) Number of network sessions available

See Also: 7A) 0Bh, 7A) 04h, INT 5Ch

**INT 2Ah Function 06h****NETBIOS, LANtastic****NETWORK PRINT-STREAM CONTROL**

Specify behavior of redirected network printer output

**Call With**

AH 06h  
 AL 01h set network printer output stream to flush print job  
 02h set network printer output stream to flush print job and start new print job  
 03h flush printer output and start new print job

**Returns**

CF set on error  
 AX error code  
 CF clear if successful

**Note:** Some LANtastic versions may not support this call.

LANtastic v4.x no longer supports this call.

**See Also:** 2F 5108h 2F 5109h 2F 1125h

**INT 2Ah Function 2001h****MS Networks or NETBIOS****UNKNOWN FUNCTION**

The purpose of this function is not known.

**Note:** This function is intercepted by DMSQserv 2.x

**INT 2Ah Functions 2002h and 2003h****Network****UNKNOWN FUNCTIONS**

The purpose of these functions is not known.

**Note:** These functions are called by MS-DOS 3.30 6.00 MPFSrv

**INT 2Ah Function 7802h****PC LAN PROG v1.31+****GET LOGGED ON USER NAME****Call With**

AX 7802h  
 ES:DI pointer to 8 byte buffer to be filled

**Returns**

AX 00h if no user logged on to Extended Services  
 AX nonzero if user logged on to Extended Services

**Note:** Some LANtastic versions may not support this call.

**INT 2Ah Function 80h****Network****BEGIN DOS CRITICAL SECTION**

This function enters the DOS kernel critical section. The region of code is being executed is not allowed to be interrupted by device drivers, to avoid race conditions. This function is used by tasking device driver.

**Called With**

AH 80h  
 AL critical-section number: 00h 0Eh  
 01h DOS kernel SHARE FILE DDMGR

02h DOS kernel SHARE FILE DDMGR

07h DOS kernel DDMGR









Print command-line arguments unless MHI1121 has been loaded.

This command-line function is documented in *MS-DOS 6.0 Programmer's Reference*.

See DD21 (ASM) in Chapter 10.

#### Format of commandline

Offset	Size	Description
00h	1	command string
01h	var	command string
N	01h	CR

### INT 2Fh Function 00h PRINT.COM UNKNOWN FUNCTION DOS 2.x only

The purpose of this function has not been determined.

Call With:

AX = 00h  
*additional arguments have unknown*

Returns:

Notes: INT 2Fh, IN 00h cause PRINT to return the number of bytes in the queue in AX.  
Values in AX other than 00h or 01h cause PRINT to return the number of bytes in the queue in AX.

See Also: N121 (AH=01)

### INT 2Fh Function 0080h PRINT.COM GIVE PRINT A TIME SLICE DOS 3.1+

Allow PRINT to execute for a while.

Call With:

AX = 0080h

Returns:

after PRINT executes

### INT 2Fh Function 0106h PRINT.COM GET PRINTER DEVICE DOS 3.3+

Determine which device PRINT is using for output if there are any bytes in the print queue.

Call With:

AX = 0106h

Returns:

- CF set if bytes in print queue
  - AX = error code 0000h = queue full
  - DSI pointer to device driver header
- CF clear if print queue empty
  - AX = 0000h

Note: documented for DOS 3.x, but not documented for prior versions (The first edition of *Undocumented DOS* incorrectly documented this as "Check if Error on Output Device".)

See Also: 21 (IOPL)

### INT 2Fh Function 0200h INSTALLATION CHECK PC LAN REDIR/REDIRFS

Determine whether the PC LAN Program redirector is installed.

## Call With

AX 0200

## Returns

AX 0000 (success)

**INT 2Fh Functions 0201h to 0204h****PC LAN REDIR/REDIRIFS****UNKNOWN FUNCTIONS**

Call With: AX 0201-0204  
 Returns: AX 0000 (success)

## Call With

AX 0201-0204h

## Returns

Notes: These functions are called by DOS 5.0+ PRINT.COM

**INT 2Fh Function 0500h****DOS 3+****CRITICAL ERROR HANDLER INSTALLATION CHECK**

## Call With

AX 0500

## Returns

AX 0000 (critical error handler installed)  
 0x0001 (not installed, call must fail)  
 0x0002 (not installed)

Notes: This function is called by COMMAND.COM

See Also: [FAT](#), [INT 2Fh](#), [MH](#), [OS](#)

**INT 2Fh Function 05h****DOS 3+****CRITICAL ERROR HANDLER EXPAND ERROR INTO STRING**

(This function is only used by error-handling)

Call With:

DOS 3+

DOS 4+

Returns: AX = error code (see [INT 2Fh](#)); CX = error message address

Notes: See [INT 2Fh](#) for details on error-handling.

Call Set if error code can't be converted

**Notes:** A check of the file `COMMAND.COM` can be done by a user who has installed the system. However, if the user is not a computer administrator of the default error message Network redirectors should use this function.

**See Also:** *INT 2fh 2f 1213*

---

### INT 2Fh Function 0600h

DOS 3+

#### ASSIGN INSTALLATION CHECK

Action: to check whether ASSIGN has been loaded

Call With

AX 0600h

Returns

AX  
 0000h OK to install  
 0100h not installed, but not OK to install  
 0200h not loaded

**Note:** ASSIGN is not a VR in DR DOS 5.0, it is internally replaced by SUBST. This was corrected in a subsequent release to the release of DR DOS 6.0.

**See Also:** *2f 52h 2f 0601h*

---

### INT 2Fh Function 0601h

DOS 3+

#### ASSIGN GET DRIVE ASSIGNMENT TABLE

Action: to get the current drive assignment table for the current ASSIGN

Call With

AX 0601h

Returns

ES: segment of ASSIGN work area and assignment table

**Note:** This table is initially set to 01h 02h 03h mapped to. This table is initially set to 01h 02h 03h.

**See Also:** *2f 0600h*

---

### INT 2Fh Function 0800h

DOS 3.2+

#### DRIVER SYS SUPPORT AVAILABILITY CHECK

Action: to check the availability of the driver system

Call With

AX 0800h

Returns

AX  
 0000h OK to install OK to install  
 0100h not installed, not OK to install  
 0200h installed

**Note:** This call is supported by DR DOS 5.0

---

### INT 2Fh Function 0801h

DOS 3.2+

#### DRIVER SYS SUPPORT ADD NEW BLOCK DEVICE

Action: to add a new logical drive alias for an existing physical drive

Call With

AX 0801h  
 DS:DI pointer to drive data table (see *2f 0803h*)

**Notes:** This function is supported by DR DOS 5.0 and DR DOS 6.0.



**—command code 00h—**

00h	BYTE	return: number of units
01h	WORD	call: pointer to DOS device helper function (see below; European MS-DOS 4.0 only)
		call: pointer past end of memory available to driver (DOS 5.0)
12h	WORD	return: address of first free byte following driver call
		call: pointer to command-line arguments
		ret: rax: pointer to KPMI array (block drivers) or 0000h (character devices)
16h	BYTE	DOS 3+ drive number for first unit of block driver (0-3)

**—European MS-DOS 4.0—**

17h	WORD	pointer to function to save registers on stack
-----	------	--

**—DOS 5+—**

17h	WORD	return: error message flag; set to 0001h for MS-DOS to display error message on initialization failure
-----	------	--

**—command codes 11h, 12h—**

00h	BYTE	reserved
-----	------	----------

**—command code 15h—**

no further fields

**—command codes 13h, 19h—**

00h	BYTE	category code
		00h unknown
		01h CCMn
		03h C CN
		05h L P n
		07h mouse (European MS-DOS 4.0)
		08h disk
		91h (SMARTER) Media Access Control driver
0Eh	BYTE	function code
		00h (SMARTER) MAC Band request
01h	WORD	copy of CS at time of IOCTL call (apparently reserved in DOS 3.3)
		SI contains European MS-DOS 4.0 reserved; DOS 5.0
11h	WORD	offset of device driver header
		DI contains European MS-DOS 4.0 reserved; DOS 5.0
13h	WORD	pointer to parameter block from 21-4401h or AX-4401h

**Values for error code:**

00h	write protect violation
01h	unknown unit
02h	drive not ready
03h	unknown command
04h	CRC error
05h	bad drive request structure length
06h	seek error
07h	unknown media
08h	sector not found
09h	printer out of paper
0Ah	write fault
0Bh	read fault
0Ch	general failure
0Dh	reserved
0Eh	C/D ROM media unavailable

00h invalid disk change

### Call European MS DOS 4.0 device helper function with:

DI = function

00h "Schedule lock" called on each timer tick

AI = tick interval in milliseconds

01h "DevDone" device I/O complete

ES:BX = pointer to request header

Note: `IOPL` register status word may be called from an interrupt handler

02h "PullRequest" pull next request from queue

ES:SI = `IOPL` pointer to start of device request queue

Return: ZF clear if pending request

ES:BX = pointer to request header

ZF set if no more requests

03h "PopParticular" remove specific request from queue

ES:SI = `IOPL` pointer to start of device request queue

ES:BX = pointer to request header

Return: ZF set if request header not found

04h "PushRequest" push the request onto the queue

ES:SI = `IOPL` pointer to start of device request queue

ES:BX = pointer to request header

interrupts disabled

05h "ControlSpillFilter" keyboard input check

AX = character (high byte 00h if PC-ANSI character)

Return: ZF set if character should be discarded

ZF clear if character should be handled normally

Note: `IOPL` register status word may be called from an interrupt handler

06h "SendSignal" send signal on keyboard event

ES:SI = `IOPL` pointer to start of device request queue

ES:BX = pointer to request header

interrupts disabled

07h "SignalEvent" send signal on keyboard event

AI = event identifier

Return: AI:HI AX destroyed

08h "ProcBlock" block on event

AX:BX = event identifier (typically a pointer)

CX = timeout in ms or 0000h for never

• `IOPL` register status word may be called from an interrupt handler

interrupts disabled

Return: after corresponding ProcRun call

CF clear if event wakeup, set if unusual wakeup

ZF set if timeout wakeup, clear if interrupted

AI = wakeup code (nonzero if unusual wakeup)

interrupts enabled

AX:BX destroyed

Note: blocks process and schedules another to run

09h "ProcRun" unblock process

AX:BX = event identifier (typically a pointer)

Return: AX = number of processes awakened

ZF set if no processes awakened

BX:CX:DX destroyed

0Ah "QueueInit" initialize: clear character queue



DS:BX — pointer to character queue structure (see below)

**Note:** the `queue_size` field must be set before calling

01h "QueueWrite" put a character in the queue

DS:BX — pointer to character queue (see below)

AL — character to append to end of queue

**Return:** ZF set if queue is full

CF clear if character stored

02h "QueueRead" get a character from the queue

DS:BX — pointer to character queue (see below)

**Return:** ZF set if queue is empty

ZF clear if characters in queue

AL — first character in queue

10h "GetDOSV" — return pointer to DOS variable

AL — index of variable

B8h — current process ID

BX — index into variable if AL specifies an array

CX — expected length of variable

**Return:** CF clear if successful

DS:AX — pointer to variable

CF set on error

AX,DX destroyed

BX,CX destroyed

**Note:** the variables may not be modified

14h "Yield" yield CPU if higher priority task ready to run

**Return:** ZF MS-DOS error

1Bh "CriticalSectionBegin" begin system critical section

DS:BX — pointer to semaphore (6 BYTES, initialized to zero)

**Return:** AX,BX,CX,DX destroyed

1Ch "CriticalSectionEnd" end system critical section

DS:BX — pointer to semaphore (6 BYTES, initialized to zero)

**Return:** AX,BX,CX,DX destroyed

**Note:** must be called by the owner of the process which called `CriticalEnter` on the semaphore

**Note:** The DWORD pointing at the request cache must be allocated by the driver and initialized to 0000h 0000h — it always points at the next request to be executed

#### Format of character queue:

Offset	Size	Description
00h	WORD	size of queue in bytes
02h	WORD	index of next character out
04h	WORD	count of characters in the queue
06h	N BYTES	queue buffer

### INT 2Fh Function 0803h

DOS 4+

#### DRIVER.SYS support GET DRIVE DATA TABLE LIST

Return a pointer to the first in a list of drive data tables describing the layout of the logical drives supported by the combination of the default disk device driver and aliases established with DRIVER.SYS

Call With

AX 0803h

## Returns

DWORD pointer to first drive data table in list

Note This function is not available under IBM DOS 5.0

See Also: 21 0801h

## Format of DOS 3.30 drive data table:

Offset	Size	Description
00h	WORD	pointer to next table
01h	BYTE	physical unit number (for INT 13h)
05h	BYTE	logical drive number (0-5)
06h	19 BYTES	BIOS Parameter Block; see also 21 53h
Offset	Size	Description
00h	WORD	bytes per sector
02h	BYTE	sectors per cluster, 1 if of unknown
03h	WORD	number of reserved sectors
05h	BYTE	number of FATs
06h	WORD	number of root dir entries
08h	WORD	total sectors
0Ah	BYTE	media descriptor (00h if unknown)
0Bh	WORD	sectors per FAT
0Dh	WORD	sectors per track
0Fh	WORD	number of heads
11h	WORD	number of hidden sectors
19h	BYTE	<ul style="list-style-type: none"> <li>0 FAT-12 instead of 12 bit FAT</li> <li>1 512 byte sectors</li> <li>2 1024 byte sectors</li> <li>3 2048 byte sectors</li> <li>4 4096 byte sectors</li> <li>5 8192 byte sectors</li> <li>6 16384 byte sectors</li> <li>7 32768 byte sectors</li> <li>8 65536 byte sectors</li> <li>9 131072 byte sectors</li> <li>A 262144 byte sectors</li> <li>B 524288 byte sectors</li> <li>C 1048576 byte sectors</li> <li>D 2097152 byte sectors</li> <li>E 4194304 byte sectors</li> <li>F 8388608 byte sectors</li> </ul>
1A	WORD	<ul style="list-style-type: none"> <li>0 512 byte sectors</li> <li>1 1024 byte sectors</li> <li>2 2048 byte sectors</li> <li>3 4096 byte sectors</li> <li>4 8192 byte sectors</li> <li>5 16384 byte sectors</li> <li>6 32768 byte sectors</li> <li>7 65536 byte sectors</li> <li>8 131072 byte sectors</li> <li>9 262144 byte sectors</li> <li>A 524288 byte sectors</li> <li>B 1048576 byte sectors</li> <li>C 2097152 byte sectors</li> <li>D 4194304 byte sectors</li> <li>E 8388608 byte sectors</li> <li>F 16777216 byte sectors</li> </ul>
1B	WORD	<ul style="list-style-type: none"> <li>0 512 byte sectors</li> <li>1 1024 byte sectors</li> <li>2 2048 byte sectors</li> <li>3 4096 byte sectors</li> <li>4 8192 byte sectors</li> <li>5 16384 byte sectors</li> <li>6 32768 byte sectors</li> <li>7 65536 byte sectors</li> <li>8 131072 byte sectors</li> <li>9 262144 byte sectors</li> <li>A 524288 byte sectors</li> <li>B 1048576 byte sectors</li> <li>C 2097152 byte sectors</li> <li>D 4194304 byte sectors</li> <li>E 8388608 byte sectors</li> <li>F 16777216 byte sectors</li> </ul>
1C	WORD	<ul style="list-style-type: none"> <li>0 512 byte sectors</li> <li>1 1024 byte sectors</li> <li>2 2048 byte sectors</li> <li>3 4096 byte sectors</li> <li>4 8192 byte sectors</li> <li>5 16384 byte sectors</li> <li>6 32768 byte sectors</li> <li>7 65536 byte sectors</li> <li>8 131072 byte sectors</li> <li>9 262144 byte sectors</li> <li>A 524288 byte sectors</li> <li>B 1048576 byte sectors</li> <li>C 2097152 byte sectors</li> <li>D 4194304 byte sectors</li> <li>E 8388608 byte sectors</li> <li>F 16777216 byte sectors</li> </ul>
1D	WORD	<ul style="list-style-type: none"> <li>0 512 byte sectors</li> <li>1 1024 byte sectors</li> <li>2 2048 byte sectors</li> <li>3 4096 byte sectors</li> <li>4 8192 byte sectors</li> <li>5 16384 byte sectors</li> <li>6 32768 byte sectors</li> <li>7 65536 byte sectors</li> <li>8 131072 byte sectors</li> <li>9 262144 byte sectors</li> <li>A 524288 byte sectors</li> <li>B 1048576 byte sectors</li> <li>C 2097152 byte sectors</li> <li>D 4194304 byte sectors</li> <li>E 8388608 byte sectors</li> <li>F 16777216 byte sectors</li> </ul>
1E	WORD	<ul style="list-style-type: none"> <li>0 512 byte sectors</li> <li>1 1024 byte sectors</li> <li>2 2048 byte sectors</li> <li>3 4096 byte sectors</li> <li>4 8192 byte sectors</li> <li>5 16384 byte sectors</li> <li>6 32768 byte sectors</li> <li>7 65536 byte sectors</li> <li>8 131072 byte sectors</li> <li>9 262144 byte sectors</li> <li>A 524288 byte sectors</li> <li>B 1048576 byte sectors</li> <li>C 2097152 byte sectors</li> <li>D 4194304 byte sectors</li> <li>E 8388608 byte sectors</li> <li>F 16777216 byte sectors</li> </ul>
1F	WORD	<ul style="list-style-type: none"> <li>0 512 byte sectors</li> <li>1 1024 byte sectors</li> <li>2 2048 byte sectors</li> <li>3 4096 byte sectors</li> <li>4 8192 byte sectors</li> <li>5 16384 byte sectors</li> <li>6 32768 byte sectors</li> <li>7 65536 byte sectors</li> <li>8 131072 byte sectors</li> <li>9 262144 byte sectors</li> <li>A 524288 byte sectors</li> <li>B 1048576 byte sectors</li> <li>C 2097152 byte sectors</li> <li>D 4194304 byte sectors</li> <li>E 8388608 byte sectors</li> <li>F 16777216 byte sectors</li> </ul>
20	WORD	<ul style="list-style-type: none"> <li>0 512 byte sectors</li> <li>1 1024 byte sectors</li> <li>2 2048 byte sectors</li> <li>3 4096 byte sectors</li> <li>4 8192 byte sectors</li> <li>5 16384 byte sectors</li> <li>6 32768 byte sectors</li> <li>7 65536 byte sectors</li> <li>8 131072 byte sectors</li> <li>9 262144 byte sectors</li> <li>A 524288 byte sectors</li> <li>B 1048576 byte sectors</li> <li>C 2097152 byte sectors</li> <li>D 4194304 byte sectors</li> <li>E 8388608 byte sectors</li> <li>F 16777216 byte sectors</li> </ul>
21	WORD	<ul style="list-style-type: none"> <li>0 512 byte sectors</li> <li>1 1024 byte sectors</li> <li>2 2048 byte sectors</li> <li>3 4096 byte sectors</li> <li>4 8192 byte sectors</li> <li>5 16384 byte sectors</li> <li>6 32768 byte sectors</li> <li>7 65536 byte sectors</li> <li>8 131072 byte sectors</li> <li>9 262144 byte sectors</li> <li>A 524288 byte sectors</li> <li>B 1048576 byte sectors</li> <li>C 2097152 byte sectors</li> <li>D 4194304 byte sectors</li> <li>E 8388608 byte sectors</li> <li>F 16777216 byte sectors</li> </ul>
22	WORD	<ul style="list-style-type: none"> <li>0 512 byte sectors</li> <li>1 1024 byte sectors</li> <li>2 2048 byte sectors</li> <li>3 4096 byte sectors</li> <li>4 8192 byte sectors</li> <li>5 16384 byte sectors</li> <li>6 32768 byte sectors</li> <li>7 65536 byte sectors</li> <li>8 131072 byte sectors</li> <li>9 262144 byte sectors</li> <li>A 524288 byte sectors</li> <li>B 1048576 byte sectors</li> <li>C 2097152 byte sectors</li> <li>D 4194304 byte sectors</li> <li>E 8388608 byte sectors</li> <li>F 16777216 byte sectors</li> </ul>
23	WORD	<ul style="list-style-type: none"> <li>0 512 byte sectors</li> <li>1 1024 byte sectors</li> <li>2 2048 byte sectors</li> <li>3 4096 byte sectors</li> <li>4 8192 byte sectors</li> <li>5 16384 byte sectors</li> <li>6 32768 byte sectors</li> <li>7 65536 byte sectors</li> <li>8 131072 byte sectors</li> <li>9 262144 byte sectors</li> <li>A 524288 byte sectors</li> <li>B 1048576 byte sectors</li> <li>C 2097152 byte sectors</li> <li>D 4194304 byte sectors</li> <li>E 8388608 byte sectors</li> <li>F 16777216 byte sectors</li> </ul>
24	WORD	<ul style="list-style-type: none"> <li>0 512 byte sectors</li> <li>1 1024 byte sectors</li> <li>2 2048 byte sectors</li> <li>3 4096 byte sectors</li> <li>4 8192 byte sectors</li> <li>5 16384 byte sectors</li> <li>6 32768 byte sectors</li> <li>7 65536 byte sectors</li> <li>8 131072 byte sectors</li> <li>9 262144 byte sectors</li> <li>A 524288 byte sectors</li> <li>B 1048576 byte sectors</li> <li>C 2097152 byte sectors</li> <li>D 4194304 byte sectors</li> <li>E 8388608 byte sectors</li> <li>F 16777216 byte sectors</li> </ul>
25	WORD	<ul style="list-style-type: none"> <li>0 512 byte sectors</li> <li>1 1024 byte sectors</li> <li>2 2048 byte sectors</li> <li>3 4096 byte sectors</li> <li>4 8192 byte sectors</li> <li>5 16384 byte sectors</li> <li>6 32768 byte sectors</li> <li>7 65536 byte sectors</li> <li>8 131072 byte sectors</li> <li>9 262144 byte sectors</li> <li>A 524288 byte sectors</li> <li>B 1048576 byte sectors</li> <li>C 2097152 byte sectors</li> <li>D 4194304 byte sectors</li> <li>E 8388608 byte sectors</li> <li>F 16777216 byte sectors</li> </ul>
26	WORD	<ul style="list-style-type: none"> <li>0 512 byte sectors</li> <li>1 1024 byte sectors</li> <li>2 2048 byte sectors</li> <li>3 4096 byte sectors</li> <li>4 8192 byte sectors</li> <li>5 16384 byte sectors</li> <li>6 32768 byte sectors</li> <li>7 65536 byte sectors</li> <li>8 131072 byte sectors</li> <li>9 262144 byte sectors</li> <li>A 524288 byte sectors</li> <li>B 1048576 byte sectors</li> <li>C 2097152 byte sectors</li> <li>D 4194304 byte sectors</li> <li>E 8388608 byte sectors</li> <li>F 16777216 byte sectors</li> </ul>
27	BYTE	termination null for volume label
28h	BYTE	device type (see 21 4401h)
29h	WORD	bit flags describing drive <ul style="list-style-type: none"> <li>0 fixed media</li> <li>1 slow seek supported</li> <li>2 unknown used in determining BPH to set for 21 4401h</li> <li>3 no sectors in a track are the same sex</li> <li>4 physical drive has multiple logical units</li> <li>5 current logical drive not physical drive</li> </ul>
2Bh	WORD	<ul style="list-style-type: none"> <li>0 related to disk status detection</li> <li>1 512 byte sectors</li> <li>2 1024 byte sectors</li> <li>3 2048 byte sectors</li> <li>4 4096 byte sectors</li> <li>5 8192 byte sectors</li> <li>6 16384 byte sectors</li> <li>7 32768 byte sectors</li> <li>8 65536 byte sectors</li> <li>9 131072 byte sectors</li> <li>A 262144 byte sectors</li> <li>B 524288 byte sectors</li> <li>C 1048576 byte sectors</li> <li>D 2097152 byte sectors</li> <li>E 4194304 byte sectors</li> <li>F 8388608 byte sectors</li> </ul>
2Ch	19 BYTES	BIOS Parameter Block; see also 21 53h
40h	4 BYTES	unknown
43h	4 BYTES	file name; default is "NO NAME"
44h	4 BYTES	<ul style="list-style-type: none"> <li>0 512 byte sectors</li> <li>1 1024 byte sectors</li> <li>2 2048 byte sectors</li> <li>3 4096 byte sectors</li> <li>4 8192 byte sectors</li> <li>5 16384 byte sectors</li> <li>6 32768 byte sectors</li> <li>7 65536 byte sectors</li> <li>8 131072 byte sectors</li> <li>9 262144 byte sectors</li> <li>A 524288 byte sectors</li> <li>B 1048576 byte sectors</li> <li>C 2097152 byte sectors</li> <li>D 4194304 byte sectors</li> <li>E 8388608 byte sectors</li> <li>F 16777216 byte sectors</li> </ul>
4F	16h	<ul style="list-style-type: none"> <li>0 512 byte sectors</li> <li>1 1024 byte sectors</li> <li>2 2048 byte sectors</li> <li>3 4096 byte sectors</li> <li>4 8192 byte sectors</li> <li>5 16384 byte sectors</li> <li>6 32768 byte sectors</li> <li>7 65536 byte sectors</li> <li>8 131072 byte sectors</li> <li>9 262144 byte sectors</li> <li>A 524288 byte sectors</li> <li>B 1048576 byte sectors</li> <li>C 2097152 byte sectors</li> <li>D 4194304 byte sectors</li> <li>E 8388608 byte sectors</li> <li>F 16777216 byte sectors</li> </ul>
<b>—removable media—</b>		
41h	16 WORDS	<ul style="list-style-type: none"> <li>0 512 byte sectors</li> <li>1 1024 byte sectors</li> <li>2 2048 byte sectors</li> <li>3 4096 byte sectors</li> <li>4 8192 byte sectors</li> <li>5 16384 byte sectors</li> <li>6 32768 byte sectors</li> <li>7 65536 byte sectors</li> <li>8 131072 byte sectors</li> <li>9 262144 byte sectors</li> <li>A 524288 byte sectors</li> <li>B 1048576 byte sectors</li> <li>C 2097152 byte sectors</li> <li>D 4194304 byte sectors</li> <li>E 8388608 byte sectors</li> <li>F 16777216 byte sectors</li> </ul>
<b>fixed media—</b>		
41h	WORD	<ul style="list-style-type: none"> <li>0 512 byte sectors</li> <li>1 1024 byte sectors</li> <li>2 2048 byte sectors</li> <li>3 4096 byte sectors</li> <li>4 8192 byte sectors</li> <li>5 16384 byte sectors</li> <li>6 32768 byte sectors</li> <li>7 65536 byte sectors</li> <li>8 131072 byte sectors</li> <li>9 262144 byte sectors</li> <li>A 524288 byte sectors</li> <li>B 1048576 byte sectors</li> <li>C 2097152 byte sectors</li> <li>D 4194304 byte sectors</li> <li>E 8388608 byte sectors</li> <li>F 16777216 byte sectors</li> </ul>

4Hh WORD absolute cylinder number of partition's start on physical drive (always FFFFh if primary partition)

### Format of COMPAQ DOS 3.31 drive data table:

Offset	Size	Description
00h	DWORD	pointer to next table
04h	BYTE	physical drive number (for INT 13h)
05h	BYTE	logical drive number (0-A)
06h	25 BYTES	BIOS Parameter Block - see DOS 4.01 drive data table below
17h	6 BYTES	apparently always zero
25h	BYTE	flags bit 6-16 bit F-A1 instead of 12 bit F-A1 5 large volume 1500 cylinders
26	WORD	volume label or "NO NAME" if none, always "NO NAME" for fixed media
28h	11 BYTES	terminating null for volume label
33h	BYTE	device type, see 21-440Dh
34h	BYTE	bit flags describing drive
35h	WORD	number of cylinders
37h	WORD	BIOS parameter block for highest capacity drive supported
39h	25 BYTES	apparently always zero
52h	6 BYTES	least significant byte of last accessed cylinder number
58h	BYTE	remains of last accessed cylinder number
<b>—removable media—</b>		
59h	DWORD	time of last access in clock ticks, FFFFFFFFh if never
<b>—fixed media—</b>		
59h	WORD	partition, FFFFh if primary, 0000h if extended
5Bh	WORD	absolute cylinder number of partition's start on physical drive (always FFFFh if primary partition)

### Format of DOS 4.0-6.0 drive data table:

Offset	Size	Description
00h	DWORD	pointer to next table
04h	BYTE	physical drive number (for INT 13h)
05h	BYTE	logical drive number (0-A)
06h	25 BYTES	BIOS Parameter Block - see also 21-54h
Offset	Size	Description
00h	WORD	bytes per sector
02h	BYTE	sectors per cluster, 11h if link tower
03h	WORD	number of reserved sectors
05h	BYTE	number of F-A's
06h	WORD	number of root dir entries
08h	WORD	total sectors - see offset 15h if zero
0A	BYTE	media descriptor, 00h if unknown
08h	WORD	sectors per F-A1
0Dh	WORD	sectors per track
0Fh	WORD	number of heads
11h	DWORD	number of hidden sectors
15h	DWORD	total sectors if WORD at 08h is zero
19h	BYTE	file system flags bit 6-16 bit F-A1 instead of 12 bit
20h	2 BYTES	bit 7 - unsupported disk, all accesses will return Not Ready device-open count



**INT 2Fh Function 1080h****DOS 4.x only****SHARE TURN ON FILE SHARING CHECKS**

Enables access to existing modes that has already been opened. Also can be used to open a file in share mode.

Call With

AX 1080h

Returns

AL status  
 10h successful  
 11h checking was already on

See Also: 2E 1000h 2E 1081h

**INT 2Fh Function 1081h****DOS 4.x only****SHARE TURN OFF FILE SHARING CHECKS**

Disables access to existing modes that has already been opened. Also can be used to open a file in share mode.

Call With

AX 1081h

Returns

AL status  
 10h successful  
 11h checking was already off

See Also: 2E 1000h 2E 1080h

**INT 2Fh Function 1100h****DOS 3.1+****NETWORK REDIRECTOR INSTALLATION CHECK**

Determines whether network redirector is installed. DOS 3.1+ only. (This is not a documented function.)

Call With

AX 1100h

Returns

AL 00h not installed, OK to install  
 01h not installed, OK to install  
 11h installed

Notes: This function is called by the 386 3.1+ kernel.

1. This function is not documented in the DOS 3.1+ kernel source code.  
 2. This function is not documented in the DOS 3.1+ kernel source code.  
 3. This function is not documented in the DOS 3.1+ kernel source code.  
 4. This function is not documented in the DOS 3.1+ kernel source code.  
 5. This function is not documented in the DOS 3.1+ kernel source code.  
 6. This function is not documented in the DOS 3.1+ kernel source code.  
 7. This function is not documented in the DOS 3.1+ kernel source code.  
 8. This function is not documented in the DOS 3.1+ kernel source code.  
 9. This function is not documented in the DOS 3.1+ kernel source code.  
 10. This function is not documented in the DOS 3.1+ kernel source code.  
 11. This function is not documented in the DOS 3.1+ kernel source code.  
 12. This function is not documented in the DOS 3.1+ kernel source code.  
 13. This function is not documented in the DOS 3.1+ kernel source code.  
 14. This function is not documented in the DOS 3.1+ kernel source code.  
 15. This function is not documented in the DOS 3.1+ kernel source code.  
 16. This function is not documented in the DOS 3.1+ kernel source code.  
 17. This function is not documented in the DOS 3.1+ kernel source code.  
 18. This function is not documented in the DOS 3.1+ kernel source code.  
 19. This function is not documented in the DOS 3.1+ kernel source code.  
 20. This function is not documented in the DOS 3.1+ kernel source code.  
 21. This function is not documented in the DOS 3.1+ kernel source code.  
 22. This function is not documented in the DOS 3.1+ kernel source code.  
 23. This function is not documented in the DOS 3.1+ kernel source code.  
 24. This function is not documented in the DOS 3.1+ kernel source code.  
 25. This function is not documented in the DOS 3.1+ kernel source code.  
 26. This function is not documented in the DOS 3.1+ kernel source code.  
 27. This function is not documented in the DOS 3.1+ kernel source code.  
 28. This function is not documented in the DOS 3.1+ kernel source code.  
 29. This function is not documented in the DOS 3.1+ kernel source code.  
 30. This function is not documented in the DOS 3.1+ kernel source code.  
 31. This function is not documented in the DOS 3.1+ kernel source code.  
 32. This function is not documented in the DOS 3.1+ kernel source code.  
 33. This function is not documented in the DOS 3.1+ kernel source code.  
 34. This function is not documented in the DOS 3.1+ kernel source code.  
 35. This function is not documented in the DOS 3.1+ kernel source code.  
 36. This function is not documented in the DOS 3.1+ kernel source code.  
 37. This function is not documented in the DOS 3.1+ kernel source code.  
 38. This function is not documented in the DOS 3.1+ kernel source code.  
 39. This function is not documented in the DOS 3.1+ kernel source code.  
 40. This function is not documented in the DOS 3.1+ kernel source code.  
 41. This function is not documented in the DOS 3.1+ kernel source code.  
 42. This function is not documented in the DOS 3.1+ kernel source code.  
 43. This function is not documented in the DOS 3.1+ kernel source code.  
 44. This function is not documented in the DOS 3.1+ kernel source code.  
 45. This function is not documented in the DOS 3.1+ kernel source code.  
 46. This function is not documented in the DOS 3.1+ kernel source code.  
 47. This function is not documented in the DOS 3.1+ kernel source code.  
 48. This function is not documented in the DOS 3.1+ kernel source code.  
 49. This function is not documented in the DOS 3.1+ kernel source code.  
 50. This function is not documented in the DOS 3.1+ kernel source code.

**INT 2Fh Function 1100h****MSCDEX****INSTALLATION CHECK**

Determines whether MSCDEX is installed. (This is not a documented function.)

Call With

AX 1100h  
SI AX WORD DMBX

Returns

AL 00h not installed, OK to install

If not installed, not OK to install  
 If installed, OK  
 If not installed, not OK to install  
 If installed, OK  
 If not installed, not OK to install  
 If installed, OK

STACK: WORLD ADVICE

### INT 2Fh Function 1101h

DOS 3.1+

#### NETWORK REDIRECTOR REMOVE REMOTE DIRECTORY

Removes a directory on a network or installed drive, the system drive.

Called With

AX = 1101h  
 SS = DOS DS  
 SI = first directory pointer      pointer to fully qualified directory name  
 DI = DOS pointer                  pointer to current directory structure for drive with dir

Returns

CF set on error  
 AX = DOS error code (see 71-506)  
 CF cleared if successful

Note

See Also 71-1102h, 60h 71-1103h, 71-1105h

### INT 2Fh Function 1102h

DOS 4.x only

#### IFSFUNC EXE REMOVE REMOTE DIRECTORY

Removes a directory on a network or installed drive, the system drive.

Called With

AX = 1102h  
 SS = DOS DS  
 SI = first directory pointer      directory name  
 DI = DOS pointer                  pointer to current directory structure for drive with dir

Returns

CF set on error  
 AX = DOS error code (see 71-506)  
 CF cleared if successful

Note This function appears to be obsolete (see 21-1101h)

See Also 21-1101h

### INT 2Fh Function 1103h

DOS 3.1+

#### NETWORK REDIRECTOR MAKE REMOTE DIRECTORY

Creates a new directory on a network or installed drive, the system drive.

Called With

AX = 1103h  
 SS = DOS DS  
 SI = first directory pointer      pointer to fully qualified directory name  
 DI = DOS pointer                  pointer to current directory structure for drive with dir

Returns

CF set on error  
 AX = DOS error code (see 71-506)  
 CF cleared if successful

Note This function is called by the DOS 3.1+ kernel.

See Also 71-506 21-60h 21-1101h 21-1105h

**INT 2Fh Function 1104h****DOS 4.x only****IFSFUNC.EXE MAKE\_REMOTE\_DIRECTORY**

Create a new directory on a network or installable file system drive.

## Called With

AX 1104h  
 SS DOS DS (see 2F 1203h)  
 SI VFS file system pointer (points to either physical directory entry or  
 network file entry) (see 2F 1203h for details)  
 DI VFS pointer (points to entry for directory to be created)

## Returns

CF set on error  
 AX DOS error code (see 2F 596)  
 CF clear if successful

Note: This function is implemented by `IFSFUNC.EXE` (2F 1107).

See Also: 2F 1103h

**INT 2Fh Function 1105h****DOS 3.1+****NETWORK REDIRECTOR CHDIR**

Change current directory on a network or installable file system.

## Called With

AX 1105h  
 SS DOS DS (see 2F 1203h)  
 SI VFS file system pointer (points to either physical directory entry or  
 network file entry) (see 2F 1203h for details)  
 DI VFS pointer (points to entry for directory to be changed)

## Returns

CF set on error  
 AX DOS error code (see 2F 596)  
 CF clear if successful  
 CDS updated with new path

Notes: This function is implemented by `IFSFUNC.EXE` (2F 1107). See also 2F 1103h, 2F 1104h, 2F 1106h, 2F 1107h, 2F 1108h, 2F 1109h, 2F 110A, 2F 110B, 2F 110C, 2F 110D, 2F 110E, 2F 110F, 2F 1110, 2F 1111, 2F 1112, 2F 1113, 2F 1114, 2F 1115, 2F 1116, 2F 1117, 2F 1118, 2F 1119, 2F 111A, 2F 111B, 2F 111C, 2F 111D, 2F 111E, 2F 111F, 2F 1120, 2F 1121, 2F 1122, 2F 1123, 2F 1124, 2F 1125, 2F 1126, 2F 1127, 2F 1128, 2F 1129, 2F 112A, 2F 112B, 2F 112C, 2F 112D, 2F 112E, 2F 112F, 2F 1130, 2F 1131, 2F 1132, 2F 1133, 2F 1134, 2F 1135, 2F 1136, 2F 1137, 2F 1138, 2F 1139, 2F 113A, 2F 113B, 2F 113C, 2F 113D, 2F 113E, 2F 113F, 2F 1140, 2F 1141, 2F 1142, 2F 1143, 2F 1144, 2F 1145, 2F 1146, 2F 1147, 2F 1148, 2F 1149, 2F 114A, 2F 114B, 2F 114C, 2F 114D, 2F 114E, 2F 114F, 2F 1150, 2F 1151, 2F 1152, 2F 1153, 2F 1154, 2F 1155, 2F 1156, 2F 1157, 2F 1158, 2F 1159, 2F 115A, 2F 115B, 2F 115C, 2F 115D, 2F 115E, 2F 115F, 2F 1160, 2F 1161, 2F 1162, 2F 1163, 2F 1164, 2F 1165, 2F 1166, 2F 1167, 2F 1168, 2F 1169, 2F 116A, 2F 116B, 2F 116C, 2F 116D, 2F 116E, 2F 116F, 2F 1170, 2F 1171, 2F 1172, 2F 1173, 2F 1174, 2F 1175, 2F 1176, 2F 1177, 2F 1178, 2F 1179, 2F 117A, 2F 117B, 2F 117C, 2F 117D, 2F 117E, 2F 117F, 2F 1180, 2F 1181, 2F 1182, 2F 1183, 2F 1184, 2F 1185, 2F 1186, 2F 1187, 2F 1188, 2F 1189, 2F 118A, 2F 118B, 2F 118C, 2F 118D, 2F 118E, 2F 118F, 2F 1190, 2F 1191, 2F 1192, 2F 1193, 2F 1194, 2F 1195, 2F 1196, 2F 1197, 2F 1198, 2F 1199, 2F 119A, 2F 119B, 2F 119C, 2F 119D, 2F 119E, 2F 119F, 2F 11A0, 2F 11A1, 2F 11A2, 2F 11A3, 2F 11A4, 2F 11A5, 2F 11A6, 2F 11A7, 2F 11A8, 2F 11A9, 2F 11AA, 2F 11AB, 2F 11AC, 2F 11AD, 2F 11AE, 2F 11AF, 2F 11B0, 2F 11B1, 2F 11B2, 2F 11B3, 2F 11B4, 2F 11B5, 2F 11B6, 2F 11B7, 2F 11B8, 2F 11B9, 2F 11BA, 2F 11BB, 2F 11BC, 2F 11BD, 2F 11BE, 2F 11BF, 2F 11C0, 2F 11C1, 2F 11C2, 2F 11C3, 2F 11C4, 2F 11C5, 2F 11C6, 2F 11C7, 2F 11C8, 2F 11C9, 2F 11CA, 2F 11CB, 2F 11CC, 2F 11CD, 2F 11CE, 2F 11CF, 2F 11D0, 2F 11D1, 2F 11D2, 2F 11D3, 2F 11D4, 2F 11D5, 2F 11D6, 2F 11D7, 2F 11D8, 2F 11D9, 2F 11DA, 2F 11DB, 2F 11DC, 2F 11DD, 2F 11DE, 2F 11DF, 2F 11E0, 2F 11E1, 2F 11E2, 2F 11E3, 2F 11E4, 2F 11E5, 2F 11E6, 2F 11E7, 2F 11E8, 2F 11E9, 2F 11EA, 2F 11EB, 2F 11EC, 2F 11ED, 2F 11EE, 2F 11EF, 2F 11F0, 2F 11F1, 2F 11F2, 2F 11F3, 2F 11F4, 2F 11F5, 2F 11F6, 2F 11F7, 2F 11F8, 2F 11F9, 2F 11FA, 2F 11FB, 2F 11FC, 2F 11FD, 2F 11FE, 2F 11FF.

The root directory is the root of the file system. For a network drive, the root directory is the root directory of the network drive. For an installable file system, the root directory is the root directory of the file system.

See Also: 2F 306, 2F 606, 2F 1101h, 2F 1103h

**INT 2Fh Function 1106h****DOS 3.1+****NETWORK REDIRECTOR CLOSE\_REMOTE\_FILE**

Close a file on a network or installable file system drive.

## Called With

AX 1106h  
 FS DI pointer to SEI  
 SI DI (points to SEI) (see 2F 1103h for details)

## Returns

CF set on error  
 AX DOS error code (see 2F 596)  
 CF clear if successful

Note: This function is called by the DOS 3.1+ kernel.

See Also: *71 30b 2E 1201h 2E 1227h*  
 Chapter 8  
 See Also: *71 30b 2E 1201h 2E 1227h*

**INT 2Fh Function 1107h****DOS 3.1+****NETWORK REDIRECTOR COMMIT REMOTE FILE**

Writes a file to a remote server and flushes disk buffers for a file on a network or installable file system drive.  
 Called With

AX = 1107h  
 ES:DI = pointer to SFT  
 SFT:DPB field = pointer to DPB of drive containing file

**Returns**

CF set on error  
 AX = DOS error code (see 2F 50h)  
 CF clear if successful  
 if buffers for file flushed  
 directory entry updated

**Note:** This function is called by the DOS 3.1+ kernel.

See Also: 2F 5100h 2F 08h

**INT 2Fh Function 1108h****DOS 3.1+****NETWORK REDIRECTOR READ FROM REMOTE FILE**

Reads a file from a remote server for a file on a network or installable file system drive.

**Called With**

AX = 1108h  
 CX = number of bytes  
 SS = DOS DS (see 2F 1203h)  
 SDA:DATA field = pointer to user buffer  
 ES:DI = pointer to SFT  
 SFT:DPB field = pointer to DPB of drive containing file

**Returns**

CF set on error  
 AX = DOS error code (see 2F 50h)  
 CF clear if successful  
 CX = number of bytes read (0000h if at end of file)  
 SFT updated

**Note:** If the remote server is a DOS 3.1+ server, the file is read from a local file on the server. See *EMsgs 8 30* and 8 31.

See Also: *71 00b 2E 5100h 2F 1109h 2F 1229h*

**INT 2Fh Function 1109h****DOS 3.1+****NETWORK REDIRECTOR WRITE TO REMOTE FILE**

Writes a file to a remote server for a file on a network or installable file system drive.

**Called With**

AX = 1109h  
 CX = number of bytes  
 SS = DOS DS  
 SDA:DATA field = pointer to user buffer  
 ES:DI = pointer to SFT  
 SFT:DPB field = pointer to DPB of drive containing file



**Returns**

CF set on error  
 AX DOS error code (see 2f 59b)  
 CF clear if successful  
 CX number of bytes written  
 SI updated

**Note:** This function is called by the DOS 3.1+ kernel

**See Also:** 2f 49b, 2f 5300b, 2f 1107b, 2f 1108b

**INT 2Fh Function 110Ah****DOS 3.1-3.31 only****NETWORK REDIRECTOR - LOCK REGION OF FILE**

Request that another process be allowed access to a portion of the specified file.

**Called With**

AX 110Ah  
 BX file handle  
 CX DX starting offset  
 SI high word of size  
 SS DOS DS (see 2f 1203h)  
 ES:DI pointer to SFI  
 SFI:DI:CDI pointer to DPF of drive containing file  
 SI:ACK WORD low word of size

**Returns**

CF set on error  
 AX DOS error code (see 2f 59b)  
 SI:ACK unchanged

**Notes:** This function is called by the DOS 3.10-3.31 kernel. The redirector is expected to resolve lock conflicts.

**See Also:** 2f 56b, 2f 110Bb

**INT 2Fh Function 110Ah****DOS 4+****NETWORK REDIRECTOR - LOCK/UNLOCK REGION OF FILE**

Deny or allow access to a portion of a remote file by other processes.

**Called With**

AX 110Ah  
 BL function  
 00h lock  
 01h unlock  
 DS:DX pointer to parameter block (see below)  
 SS DOS DS  
 ES:DI pointer to SFI  
 SFI:DPF:CDI pointer to DPF of drive containing file

**Returns**

CF set on error  
 AX DOS error code (see 2f 59b)

**Notes:** This function is called by the DOS 4.0+ kernel.

The redirector is expected to resolve lock conflicts.

**See Also:** 2f 56b, 2f 110Bb

**Format of parameter block**

Offset	Size	Description
00h	DWORD	Start offset
04h	DWORD	Size of region

**INT 2Fh Function 1108h****DOS 3.1 3.31 only****NETWORK REDIRECTOR UNLOCK REGION OF FILE**

Allow other processes to access the specified portion of the file.

**Called With**

AX	DWORD
JX	FILE_NAME
SI	high word of size
EDI	WORD low word
EDI	pointer to SEI for SET FILE LOCK structure

**Returns**

CF	set on error
SI	DWORD low word
SI	high word

**Notes**See Also: [File Locking](#), [File Locking](#), [File Locking](#), [File Locking](#)**INT 2Fh Function 110Ch****DOS 3.1+****NETWORK REDIRECTOR GET DISK SPACE**

Get the amount of free space on a disk.

**Called With**

AX	DWORD
EDI	pointer to DOS FAT sectors structure

**Returns**

SI	sectors
DI	bytes
SI	total
SI	bytes per sector
SI	total

Note: Only use on local drives with DOS 3.1+.

See Also: [File Locking](#)**INT 2Fh Function 110Eh****DOS 3.1+****NETWORK REDIRECTOR SET REMOTE FILE'S ATTRIBUTES**

Change the attributes of a file on a network or system block system drive.

**Called With**

AX	DWORD
SI	DWORD low
SI	high word (same pointer)
SI	pointer
SI	WORD low

**Returns**

CF set on error  
 AX DOS error code (see 27-50b)  
 CF clear if successful  
 SI/AX unchanged

**Note:** This function is called by the DOS 3.1+ kernel

See Also: 21-430h, 21-60b, 21-110h

**INT 2Fh Function 110Fh****DOS 3.1+****NETWORK REDIRECTOR - Get Remote File's Attributes And Size**

Get file attributes of a file on a network or stub file system.

**Called With**

AX 110Fh  
 SS DOS DS  
 SI file name pointer (fully qualified)  
 DI file name pointer (current directory structure for drive with file)

**Returns**

CF set on error  
 AX DOS error code (see 27-50b)  
 CF clear if successful  
 AX file attributes  
 BX/DI file size

**Note:** This function is called by the DOS 3.1+ kernel

See Also: 21-430b, 21-60b, 21-110h

**INT 2Fh Function 1111h****DOS 3.1+****NETWORK REDIRECTOR - RENAME REMOTE FILE**

Rename a file on a network or stub file system.

**Called With**

AX 1111h  
 SS DOS DS  
 SI first file name pointer (fully qualified old name)  
 DI second file name pointer (fully qualified new name)  
 DI/CX file name pointer (current directory structure for drive with file)

**Returns**

CF set on error  
 AX DOS error code (see 27-50b)  
 CF clear if successful

**Note:** This function is called by the DOS 3.1+ kernel

See Also: 21-56d, 21-60b

**INT 2Fh Function 1113h****DOS 3.1+****NETWORK REDIRECTOR - DELETE REMOTE FILE**

Remove a file from a network or installed file system drive.

**Called With**

AX 1113h  
 SS DOS DS

DS: DOS DS

DOS DS

### Returns

AX: DOS error code (see 71-50b)

CF: clear if successful

**Note:** This function is called by the DOS 3.1 kernel.

For file specification may contain wildcards.

See Also: 71-41h, 71-60h

## INT 2Fh Function 1116h

DOS 3.1+

### Network Redirector Open Existing Remote File

Prepares for access to an existing file located on a network drive.

#### Called With

AX: 1116h  
 ES:DI: pointer to uninitialized SEI  
 SS: DOS DS  
 SI: offset to WORD file open mode (see 21-70b)

#### Returns

CF: set on error  
 AX: DOS error code (see 21-50b)  
 CF: clear if successful  
 SEI field: except handle count, which DOS manages itself  
 SI: WORD file mode

**Note:** This function is called by the DOS 3.1 kernel.

See Also: 71-20b, 71-60b, 21-4100h, 21-4117h, 21-4118h, 21-4119h, 21-4120h

## INT 2Fh Function 1117h

DOS 3.1+

### Network Redirector Create/truncate Remote File With Cds

Create a file on a network drive, or truncate an existing file to zero length.

#### Called With

AX: 1117h  
 ES:DI: pointer to uninitialized SEI  
 SS: DOS DS  
 SI: offset to string pointer (pointer to fully qualified name of file to open)  
 SI+6: offset to WORD file creation mode  
 SI+8: offset to file attributes  
 SI+10: offset to WORD file mode (create, 01h; create new file, 02h)

#### Returns

CF: set on error  
 AX: DOS error code (see 21-50b)  
 CF: clear if successful  
 SEI field: except handle count, which DOS manages itself  
 SI: WORD file mode

**Note:** This function is called by the DOS 3.1 kernel.

See Also: 21-4100h, 21-4117h, 21-4118h, 21-4119h, 21-4120h, 21-4121h

**INT 2Fh Function 1118h****DOS 3.1+****Network Redirector Create/truncate File Without Cds**

Create a file on a drive which does not have a current directory structure.

**Called With:**

AX	1118h
DI:DI	pointer to uninitialized SI
SI	DOS DS
SDI first filename pointer	pointer to first filename character
SI:AK	WORD file creation mode low byte = file attributes high byte = 00h normal create, 01h create new file

**Returns:**

SI:DI:DI  
SI:AK unchanged

**Note:** This function is available in DOS 3.1 and later versions. It is not available in versions of DOS 3.0. SDI:DS:DI:DI must be set to the drive and directory to create the file. SDI:DS:DI must be set to the drive and directory to create the file. Create New file for the name: 0400:BAR.

See Also: 2F:00h, 2F:1106h, 2F:1116h, 2F:1117h, 2F:1118h

**INT 2Fh Function 1119h****DOS 3.1+****NETWORK REDIRECTOR FIND FIRST FILE WITHOUT CDS**

Begin a directory search on a network or local drive which does not have a current directory structure.

**Called With:**

AX	1119h
SI	DOS DS
DI	DOS DS
DI:DI	pointer to first filename character
SDI first filename pointer	pointer to first filename character
SDI search attribute	01h search for file

**Returns:**

CF set on error  
AX DOS error code (see 2F:800h)  
CF clear if successful  
DI:DI pointer to first file found  
DI:DI:DI:DI pointer to first byte must be set  
DI:DI:DI:DI standard directory entry for file

**Notes:** This function is available in DOS 3.1 and later versions. SDI:DS:DI:DI must be set to the drive and directory to search. DOS 4.x IFSI:NC returns CF set, AX=0003h.

**INT 2Fh Function 111Bh****DOS 3.1+****NETWORK REDIRECTOR - FINDFIRST**

Begin a directory search on a network or local drive which does not have a current directory structure.

**Called With:**

AX	111Bh
SI	DOS DS
DI	DOS DS
DI:DI	pointer to first filename character
SDI first filename pointer	pointer to first filename character
SDI search attribute	01h search for file

SI: Attribute mask for search.

#### Returns

- CF: Set on error.
- AX: DOS error code (see 2F 59h).
- DI: Offset of virus.
- DI+1: Updated (if offset search data) but "of last byte" must be set.
- DI+15h: Standard directory entry for file.

#### Notes

See Also 2F 41h, 2F 60h, 7E 111Ch.

### INT 2Fh Function 111Ch

DOS 3.1+

#### NETWORK REDIRECTOR FINDNEXT

Continue a directory search on a network or installed file system drive.

#### Called With

- AX: 111Ch.
- DS: DOS DS.
- DI: 2F into (offset search data) (see 2F 41h).

#### Returns

- CF: Set on error.
- AX: DOS error code (see 2F 59h).
- DI: Offset of virus.
- DI+1: Updated (if offset search data) but "of last byte" must be set.
- DI+15h: Standard directory entry for file.

#### Notes

See Also 2F 41h, 7E 1111h.

### INT 2Fh Function 111Dh

DOS 3.1+

#### Network Redirector Close All Remote Files For Process

#### Called With

- AX: 111Dh.
- DS: DOS DS.
- SI: DOS DS.

#### Returns

Notes: If successful, invoked by the DOS 3.1+ kernel.

The redirector closes all FDs opened by the process.

See Also 7E 8104h.

### INT 2Fh Function 111Eh

DOS 3.1+

#### NETWORK REDIRECTOR DO REDIRECTION

Various subfunctions allow control of network redirector.

#### Called With

- AX: 111Eh.
- SS: DOS DS.
- SI: ACK, WORD function to execute.

- 5f00h** get redirection mode  
 BL type: 03h printer, 04h disk  
**Returns:** BH state: 00h off, 01h on
- 5f01h** set redirection mode  
 BL type: 03h printer, 04h disk  
 BH state: 00h off, 01h on
- 5f02h** get redirection list entry  
 BX redirection list index  
 DS:SI pointer to 16 byte local device name buffer  
 ES:DI pointer to 128 byte network name buffer  
**Returns:** must set user's BX to device type, and CX to stored parameter value (using 2f1210h to get stack frame address)
- 5f03h** redirect device  
 BL device type (see 2f5f00h)  
 CX stored parameter value  
 DS:SI pointer to ASC I/ source device name  
 ES:DI pointer to destination ASC I/ network path + ASC IZ passwd
- 5f04h** cancel redirection  
 DS:SI pointer to ASC I/ device name or network path
- 5f05h** get redirection list extended entry  
 BX redirection list index  
 DS:SI pointer to buffer for ASC I/ source device name  
 ES:DI pointer to buffer for destination ASC I/ network path  
**Returns:** BH status flag  
 BL type: 03h printer, 04h disk  
 CX stored parameter value  
 BP NE1BIOS local session number
- 5f06h similar to 5f05h*

**Returns**

- CF set on error  
 AX error code (see 2f50h)  
 SI:ACK unchanged

**Note:** This function is called by the DOS 3.1 kernel on 2f5f0h (including LAN Manager calls).  
 See Also 2f5f00h, 2f5f01h, 2f5f02h, 2f5f03h, 2f5f04h, 2f5f05h, 2f5f06h, 2f5f05h, 2f5f06h

**INT 2Fh Function 111Fh****DOS 3.1+****NETWORK REDIRECTOR - PRINTER SETUP**

Subroutines allow getting or setting the network printer setup string or mask.

**Called With**

- AX 111Fh  
 SI:ACK WORD function  
 5E02h set printer setup  
 5f03h get printer setup  
 5f04h set printer mode  
 5E05h get printer mode

**Returns**

- CF set on error





**Returns**

CF = 0

**Note:** This function is called by the DOS 3.1+ kernel.**INT 2Fh Function 1123h****DOS 3.1+****NETWORK REDIRECTOR QUALIFY REMOTE FILENAME**

Convert a name (not absolute path name with any network redirections resolved).

**Called With**

AX = 1123h  
 DS:SI = pointer to ASCII filename to canonicalize  
 ES:DI = pointer to 128 byte buffer for qualified name

**Returns**

CF set if not resolved

**Note:** This function is called by the MS-DOS 3.1 kernel. It is not supported by DR-DOS and Linux. The name must consist of a character device.

DOS 3.1+ will attempt to resolve a file name over a network if a SMB server call is made. If this fails, DOS resolves the name locally.

See Also: 21 160h, 21 1221h

**INT 2Fh Function 1124h****DOS 3.1+****NETWORK REDIRECTOR PRINTER OFF**

Adjust network printer standard output to the 001 state.

**Called With**

AX = 1124h  
 SI = pointer to NT  
 SS = DOS:DS  
 add %s *replacement if any unknown*

**Returns**

CF = 0

**Note:** This function is called by the DOS 3.1+ kernel. (21 1126h is the old set).

See Also: 21 1126h

**INT 2Fh Function 1125h****DOS 3.1+****NETWORK REDIRECTOR REDIRECTED PRINTER MODE**

Set or kill the state of print streams for the network printer.

**Called With**

AX = 1125h  
 STACK = WORD instruction  
 5100h get print stream state  
 Returns  
 DI = current state  
 5100h set print stream state  
 DI = new state  
 5100h finish print job

**Returns**

CF set on error

AX = error code, see 21 50h

STACK unchanged

**Note:** This function is called by the DOS 3.1+ kernel.

See Also: 21 5100h, 21 51008h, 21 51009h

---

**INT 2Fh Function 1126h**  
**NETWORK REDIRECTOR PRINTER ON/OFF**
**DOS 3.1+**

Enable printer echo of standard output

**Called With**

AX 1126h  
 ES DI pointer to SF1 for file handle 4  
 SS DOS DS see 2F 1205h  
*additional arguments of any unknown*

**Returns**

CF set on error

**Note:** This function is only in DOS 3.1 and the previous version. If PrtNo changes state, a SILENT = 1 will do a carriage return. In SILENT = 0, a carriage return will do a carriage return.

See Also: 2F 1124h

---

**INT 21h Function 1127h**  
**NETWORK REDIRECTOR REMOTE COPY**
**Novell**

Network redirector function that copies a file from a remote server to a local destination. It is the opposite of the same function for more information, see chapter 8.

---

**INT 2Fh Function 1128h**  
**IFSFUNC EXE GENERIC IOCTL**
**DOS 4.x only**

Implements the "Generic IOCTL" call on network devices.

**Called With**

AX 1128h  
 SS DOS DS  
 CX function category  
 DS DI pointer to parameter block  
 EAX WORD 0-FFFFh  
*Additional arguments of any unknown*

**Returns**

CF set on error  
 AX DOS error code (see 2F 50h)  
 CF clear if successful

**Note:** This function is called by the DOS 4.0 kernel.

---

**INT 2Fh Function 112Ch**  
**IFSFUNC EXE UNKNOWN FUNCTION**
**DOS 4+**

The purpose of this function has not been determined.

**Called With**

AX 112Ch  
 SS DOS DS  
 SIA current SILENT pointer to SILENT  
*Additional arguments of any unknown*

**Returns**

CF set on error  
 AX DOS error code (see 2F 50h)  
 CF clear if successful

**Note:** Called by SHARE in DOS 5.0

**INT 2Fh Function 112Dh****DOS 4.x only****IFSFUNC.EXE - UNKNOWN FUNCTION**

The purpose of this function has not been determined.

**Called With**

AX	112Dh
BL	subfunction value of AI on INT 2Fh
	04h truncate open file to zero length
	FS:DI pointer to SFI for file
	<b>Returns:</b>
	CF clear
	<i>dir unknown</i>
	<b>Returns</b>
	CA <i>unknown (00h or 02h for DOS 4.0)</i>

FS:DI pointer to SFI

SS DOS DS

**Returns**

DS DOS DS

**Note:** This function is called by the DOS 4.0 kernel on 2F 5702h, 2F 5703h, and 2F 5704h. Perhaps 2F 112D supports OS/2 extended attributes. I was told by the DosGee, SetPath, EtcInfo functions.

**INT 2Fh Function 112Eh****DOS 4+****NETWORK REDIRECTOR EXTENDED OPEN/CREATE FILE**

Emulates the extended file open call 2F 06h for network and installable file system drives.

**Called With**

AX	112Eh
SS	DOS DS
DS	DOS DS
ES:DI	pointer to uninitialized SFI for file
SDA	extended file name pointer
SDA	extended file operation mode
SDA	extended file operation mode
SI:BX	WORD: file attributes, low byte = file attributes
	high byte = 00h if simultaneous open, 01h if create new file

**Returns**

CF	set on error
AX	error code
CF	clear if success
CX	result code
	01h file opened
	02h file created
	03h file replaced (truncated)

See also Extended, except, code count, which DOS manages itself.

**Note:** This function is called by the DOS 4.x kernel and by DOS S+ SHARE.**See Also:** 2F 06 00h, 2F 1115h, 2F 1116h, 2F 1117h

**INT 2Fh Function 1130h** **DOS 4.x only****IFSFUNC EXE GET IFSFUNC SEGMENT**

Return the segment of the resident IFSFUNC code.

Called With

AX = 1130h

Returns

ES = CS of resident IFSFUNC

**INT 2Fh Function 1200h** **DOS 3+****INTERNAL FUNCTIONS AVAILABILITY CHECK**

Determine whether the DOS internal services are present.

Call With

AX = 1200h

Returns

CF = 1 if any of the internal functions (INT 1E through 3B) are not installed

Note: See Chapters 6 and 8 for discussions of the DOS internal functions.

**INT 2Fh Function 1201h** **DOS 3+****CLOSE CURRENT FILE**

Close the file currently being operated on.

Call With

AX = 1201h

SS = DOS DS

SI = current SEI pointer      pointer to SEI of file to close

Returns

CF set on error

CF clear if successful

BX = interrupt

CX = new reference count of SEI

ESI = pointer to SEI for file

See Also: *File* 21-1106h; 21-1277h**INT 2Fh Function 1202h** **DOS 3+****GET INTERRUPT ADDRESS**

Return the address of the interrupt specified by the interrupt number.

Call With

AX = 1202h

SI = K = WORDS vector number

Returns

ESI = pointer to interrupt

SI = K and ramed

**INT 2Fh Function 1203h** **DOS 3+****GET DOS DATA SEGMENT**

Return the segment of the DOS data segment.

Call With

AX = 1203h

**Returns**

DS segment of IBM DOS COM MS-DOS MS-DOS data segment.

**Note** For DOS 3.X and 4.X, but *not* for DOS 5+, the kernel code segment is the same as the data segment.

Many applications use the DOS function 21h (calling 21/52) and using FS (ignoring BX). The "style" (DOS 3 or DOS 4) of the DS data segment is indicated by the byte at offset 4 of DOS DS (see 21/5100).

**INT 2Fh Function 1204h****DOS 3+****NORMALIZE PATH SEPARATOR**

Convert forward slashes into backslashes.

**Call With**

AX 1204h  
 STACK WORD character to normalize

**Returns**

AI 0 on success, nonzero on error (backslash or other separator)  
 ZF set if path separator  
 STACK unchanged

**INT 2Fh Function 1205h****DOS 3+****OUTPUT CHARACTER TO STANDARD OUTPUT**

Send a single character to the standard output of the current program.

**Call With**

AX 1205h  
 STACK WORD character to output

**Returns**

STACK unchanged

**Note** This function is used by the DOS command PROMPT.

**INT 2Fh Function 1206h****DOS 3+****INVOKE CRITICAL ERROR**

Causes INT 2Fh (page 784) to be invoked. AX contains the error code, and the stack contains the error message.

**Call With**

AX 1206h  
 DI error code  
 BP SI pointer to device driver name  
 SS DOS DS  
 STACK WORD value to be passed to INT 24h in AX

**Returns**

AI 0 3 for Abort, Retry, Ignore, Fail  
 STACK unchanged

See Also: INT 24h

**INT 2Fh Function 1207h****DOS 3+****MAKE DISK BUFFER MOST RECENTLY USED**

Move the pointer to the most recently used disk buffer to the end of the disk buffer segment, and the reverse of *etc*.



Used by network redirectors such as MMIOX.

See Also: INT 24h

---

### INT 2Fh Function 120Bh DOS 3+

#### SIGNAL SHARING VIOLATION TO USER

Produces a signal to the user if a sharing violation was made. The signal is a 20-bit error code (bits made with inheritance allowed).

Call With:

AX 120Bh  
 ES:DI pointer to system file table entry for previous open of file  
 SI:AX WORD extended error code (should be 20h = sharing violation)

Returns:

CF clear if operation should be retried  
 CF set if operation should not be retried  
 AX error code + 20h (see 21-59b)  
 SI:AX unchanged

Note: This function can only be called during a DOS function call. It should only be called if an attempt is made to open an already open file, or if a sharing rule is violated.

---

### INT 2Fh Function 120Ch DOS 3+

#### OPEN DEVICE AND SET SFT OWNER

Invokes the "device open" call on the device driver for the given MFI and then sets the owner of the last accessed PCB file to the calling process's ID.

Call With:

AX 120Ch  
 SI:AX current MFI pointer (pointer to MFI in PCB)  
 DS 1005:10h  
 SS 1005:10h

Returns:

ES, DI, AX destroyed (ES:DI may point to the MFI)

Note: Network redirectors must call this function for HDDFS, SCS, CD, and S.

---

### INT 2Fh Function 120Dh DOS 3+

#### GET DATE AND TIME

Return the current date and time in directory format.

Call With:

AX 120Dh  
 SS 1005:10h

Returns:

AX current date in packed format (see 21-570a)  
 DX current time in packed format (see 21-570b)

See Also: 21-24b, 21-203

---

### INT 2Fh Function 120Eh DOS 3+

#### MARK ALL DISK BUFFERS UNREFERENCED

Clears the "referenced" bit of all buffers that are automatically set when a buffer is read or written. It is set by the file system after a placement algorithm. Unreferenced buffers are generally replaced before referenced buffers.

**Call With**

AX 1201h  
 SS 1308:13h

**Returns**

SI DI pointer to first disk buffer

**See Also**

AX 1224h SEARCHING FOR A FILE

See Also 7E-1209h 2E-1210h

**INT 2Fh Function 120Fh****DOS 3+****MAKE BUFFER MOST RECENTLY USED****Call With**

AX 120Fh  
 DS DI pointer to disk buffer  
 SS 1308:13h

**Returns**

SI DI pointer to next buffer in buffer list

**See Also**

7E-1207h 7E-1208h 7E-1209h 7E-1210h 7E-1211h 7E-1212h 7E-1213h 7E-1214h 7E-1215h 7E-1216h 7E-1217h 7E-1218h 7E-1219h 7E-1220h 7E-1221h 7E-1222h 7E-1223h 7E-1224h 7E-1225h 7E-1226h 7E-1227h 7E-1228h 7E-1229h 7E-1230h 7E-1231h 7E-1232h 7E-1233h 7E-1234h 7E-1235h 7E-1236h 7E-1237h 7E-1238h 7E-1239h 7E-1240h 7E-1241h 7E-1242h 7E-1243h 7E-1244h 7E-1245h 7E-1246h 7E-1247h 7E-1248h 7E-1249h 7E-1250h 7E-1251h 7E-1252h 7E-1253h 7E-1254h 7E-1255h 7E-1256h 7E-1257h 7E-1258h 7E-1259h 7E-1260h 7E-1261h 7E-1262h 7E-1263h 7E-1264h 7E-1265h 7E-1266h 7E-1267h 7E-1268h 7E-1269h 7E-1270h 7E-1271h 7E-1272h 7E-1273h 7E-1274h 7E-1275h 7E-1276h 7E-1277h 7E-1278h 7E-1279h 7E-1280h 7E-1281h 7E-1282h 7E-1283h 7E-1284h 7E-1285h 7E-1286h 7E-1287h 7E-1288h 7E-1289h 7E-1290h 7E-1291h 7E-1292h 7E-1293h 7E-1294h 7E-1295h 7E-1296h 7E-1297h 7E-1298h 7E-1299h

If buffer or chain of it is marked unused.

See Also 2E-1207h

**INT 2Fh Function 1210h****DOS 3+****FIND UNREFERENCED DISK BUFFER****Call With**

SI DI pointer to first disk buffer to check

**Call With**

AX 1210h  
 DS DI pointer to first disk buffer to check

**Returns**

ZF clear if found  
 DS DI pointer to first unreferenced disk buffer  
 ZF set if not found

**See Also**

AX 1223h SEARCHING FOR A FILE

See Also 7E-120Fh

**INT 2Fh Function 1211h****DOS 3+****NORMALIZE ASCIZ FILENAME****Call With**

SI DI pointer to buffer containing filename

**Call With**

DI DI pointer to buffer to store file name  
 DI DI pointer to buffer for normalized filename

**Returns**

SI DI buffer index

**See Also**

INT 2Fh 2E-1221h



---

**INT 2Fh Function 1212h**  
**GET LENGTH OF ASCII STRING**
**DOS 3+**

Returns the length of a null-terminated character string.

**Call With**

AX 1212h  
 ES:DI pointer to ASCII string

**Returns**

CX length of string

See Also: 1-1275

---

**INT 2Fh Function 1213h**  
**UPPERCASE CHARACTER**
**DOS 3+**

Returns the uppercase ASCII equivalent of the lowercase ASCII character. The given character.

**Call With**

AX 1213h  
 STACK WORD character to convert to uppercase

**Returns**

AL uppercase character

STACK WORD character

---

**INT 2Fh Function 1214h**  
**COMPARE FAR POINTERS**
**DOS 3+**

Compares two 32-bit far pointers. Returns 0 if equal.

**Call With**

AX 1214h  
 DS:SI first pointer  
 ES:DI second pointer

**Returns**

ZF set if pointers are equal, ZF clear if not equal

---

**INT 2Fh Function 1215h**  
**FLUSH BUFFER**
**DOS 3+**

Flushes the disk buffer for the specified disk drive. The disk drive is identified

**Call With**

AX 1215h  
 DS:DI pointer to disk buffer  
 SS DOS DS  
 STACK WORD drives for which to skip buffer

Each drive is checked to see if the buffer is full. If so, the buffer is flushed. If the buffer is not full, the buffer is not flushed. If the buffer is not full, the buffer is not flushed.

**Returns**

STACK unchanged

**Note:** This function can only be called from within a DOS function call.

See Also: 21-1209b

**INT 2Fh Function 1216h****DOS 3+****GET ADDRESS OF SYSTEM FILE TABLE ENTRY**C. S. & S. S. *System File Table* 1216h**Call With**

AX 12,6

BX system file table entry number, such as sector id from 2E 1220

**Returns**

CF clear if success

ESI pointer to system file table entry

CX set if BX greater than 141FN

**Note:** For the implementation of this function, see Figure 6-22.

See Also: 2E 1270h

**INT 2Fh Function 1217h****DOS 3+****GET CURRENT DIRECTORY STRUCTURE FOR DRIVE****Call With**

AX 12,7

BX DOS DS

SI:BX WORD type of A:Z B:ch

**Returns**CF clear if success  
drive = (SI:BX)/0x10

CF clear if successful

DS:SI pointer to current directory structure for specified drive

SI:BX unchanged

**Note:** For implementation, see Figure 6-1.

See Also: 2E 1219

**INT 2Fh Function 1218h****DOS 3+****GET CALLER'S REGISTERS**

Returns a pointer to the stack frame containing the INT 2Fh caller's registers.

**Call With**

AX 12,8

**Returns**

DS:SI pointer to caller's AX BX CX DX SI DI BP DS:ES on stack

**Note:** The return DOS function is only valid while within a DOS function. See Figure 6-19.**INT 2Fh Function 1219h****DOS 3+****SET DRIVE****Call With**

AX 12,9

BX DOS DS

SI:BX WORD drive of default directory

**Returns**

SI:BX unchanged

**Note:** For implementation, see Figure 6-1. For the implementation of the default directory structure if made a server call, 2E 5000h.

See Also: 2E 1217h 2E 121Fh

---

**INT 2Fh Function 121Ah**  
**GET FILE'S DRIVE**
**DOS 3+**

Determine which drive a filename specifies.

**Call With**

AX 121Ah  
 DS:SI pointer to filename

**Returns**

AL drive (0=default, 1=A, etc., FF=invalid)  
 DS:SI pointer to filename with 1 leading 0, if present

**Note:** Increment SI by 2 if name starts with drive letter followed by colon.

See Also: 21-600h

---

**INT 2Fh Function 121Bh**  
**SET YEAR/LENGTH OF FEBRUARY**
**DOS 3+**

Specify the current year and return the length of February of the year, storing that length in a byte.

**Call With**

AX 121Bh  
 CX year (1980)  
 DS DOS data segment

**Returns**

AL number of days in February

See Also: 21-280h

---

**INT 2Fh Function 121Ch**  
**CHECKSUM MEMORY**
**DOS 3+**

Compute a checksum of the given range of memory. This function is also used by DOS to determine the day count since 1/1/1980 given a date.

**Call With**

AX 121Ch  
 DS:SI pointer to start of memory to checksum  
 CX number of bytes  
 AX zero if checksum  
 SS DOS DS

**Returns**

AX, CX destroyed  
 DX checksum  
 DS:SI pointer to first byte after checksum word range

See Also: 21-121Dh

---

**INT 2Fh Function 121Dh**  
**SUM MEMORY**
**DOS 3+**

Add up the values of a range of bytes and the specified limit, overflowed, and return the value which causes the limit to be exceeded. This function is also used by DOS to determine the year and month given a day count since 1/1/1980.

**Call With**

AX 121Dh  
 DS:SI pointer to memory to add up  
 CX 0000h

DX limit

**Returns**

AX byte which exceeded limit  
 CX number of bytes before limit exceeded  
 DX remainder after adding first CX bytes  
 DS:SI pointer to byte beyond the one which exceeded the

See Also: 211, 211c

**INT 2Fh Function 121Eh**

**DOS 3+**

**COMPARE FILENAMES**

**Call With**

AX 121Eh  
 DS:SI pointer to first ASCII filename  
 ES:DI pointer to second ASCII filename

**Returns**

ZF set if filenames equisvalent; ZF clear if not

See Also: 211, 211b, 211, 1221b

**INT 2Fh Function 121Fh**

**DOS 3+**

**BUILD CURRENT DIRECTORY STRUCTURE**

storage in directory was built

**Call With**

AX 121Fh  
 SS DS:DS  
 STACK WORDS, directory

**Returns**

AX WORDS, directory  
 STACK WORDS, directory

**INT 2Fh Function 1220h**

**DOS 3+**

**GET JOB FILE TABLE ENTRY**

**Call With**

AX 1220h  
 CS:SI:DI

**Returns**

CI  
 AX offset to start of file handle  
 CI clear if successful  
 ES:DI pointer to file entry for file handle in current process

**Note** The byte pointed at by ES:DI is the number of the SFI entry for the file handle or 11h if the handle exists open.

For the implementation of this function, see Figure B-22.

See Also: 211, 211c, 1221b

### INT 2Fh Function 1221h CANONICALIZE FILE NAME

DOS 3+

Given a file specification, return an absolute pathname which takes into account all renaming due to JOIN, SUBST, ASSOC, or network redirections.

#### Call With

AX	1221h
DS:SI	pointer to file name to be fully qualified
ES:DI	pointer to 128-byte buffer for result (if canonical file name)
SS	IOPLDS

#### Returns,

see 21/60h

**Note** This function can only be called from within a DOS function call, and is otherwise identical to 21/60h.

See Also: 21/60h, 2f/1123h

### INT 2Fh Function 1222h SET EXTENDED ERROR INFO

DOS 3+

Given a set of existing records, set the error class, locus, and suggested action corresponding to the current extended error code.

#### Call With

AX	1222h
SS	IOPLDS data segment
SI	pointer to 4-byte records
DI	error code, FFh = last record
SI	error class, FFh = don't change
DI	suggested action, FFh = don't change
SI	error locus, FFh = don't change

NOA error code set

#### Returns

SI destroyed

NOA error class, error locus, and suggested action fields set

**Note**, This function can only be called from within a DOS function call.

See Also: 2f/90b, 2f/1221h

### INT 2Fh Function 1223h CHECK IF CHARACTER DEVICE

DOS 3+

Determine whether the given name is the name of a character device.

#### Call With

AX	1223h	
DS	IOPLDS data segment	
SS	IOPLDS data segment	
SI(2:18h)	IOPLDS 3:10-3:30	right character name, padded name
SI(3:22h)	IOPLDS 4:0-5:0	right character name, padded name

#### Returns

CF set if no character device by that name found

CF clear if found

BH low byte of device attribute word

**Note** This function can only be called from within a DOS function call.

See Also: 2f/5100h, 2f/5101h

---

**INT 2Fh Function 1224h**  
**SHARING DELAY**
**DOS 3+****Call With**

AX 1224h  
 BX DOS DS

**Returns**

CF flag set by 21 440B (unknown server call) 21 8109

See Also [440B](#), [21 821](#)

---

**INT 2Fh Function 1225h**  
**GET LENGTH OF ASCIZ STRING**
**DOS 3+**

Returns the length of a null-terminated character string.

**Call With**

AX 1225h  
 DS:SI pointer to ASCIZ string

**Returns**

CX length of string

See Also [21 4217h](#)

---

**INT 2Fh Function 1226h**  
**OPEN FILE**
**DOS 3.3+****Call With**

AX 1226h  
 CX access mode M from 21 319  
 DS:DX pointer to ASCIZ file  
 SI DOS file number

**Returns**

CF flag set by 21 319

AX file handle

See Also [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#)

See Also [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#)

See Also [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#)

See Also [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#)

See Also [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#)

See Also [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#)

See Also [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#)

See Also [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#)

See Also [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#)

See Also [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#)

See Also [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#)

See Also [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#)

See Also [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#)

See Also [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#)

See Also [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#), [21 319](#)

---

**INT 2Fh Function 1227h**  
**CLOSE FILE**
**DOS 3.3+**

Closes a previously opened file given its handle.

**Call With**

AX file handle  
 SI DOS data segment

**Returns**

CF set on error

AI 06h invalid file handle

CF clear if successful

**Notes:** This function can only be called from within a DOS function call (i.e., otherwise equivalent to 21/3Fh).

This function is used by NLSFUNC to access COUNTRY SYS when invoked by the DOS kernel.

See Also: 21/09h, 21/106h, 21/1201, 21/1226h

**INT 2Fh Function 1228h****DOS 3.3+****MOVE FILE POINTER**

Set the current position in the given file.

**Call With**

AX 1228h

BP 4200h, 4201h, 4202h (see 2F/42h)

BX file handle

CX/DX offset in bytes

SS DOS DS

**Returns**

as for 2F/42h

**Notes:** This function is equivalent to 2F/42h. Furthermore, it can be called from outside a DOS function call.

The user stack frame pointer is set to 4200h by the BIOS and is changed into AX. COUNTRY SYS is performed. Finally, the frame pointer is restored.

This function is used by NLSFUNC to access COUNTRY SYS when invoked by the DOS kernel.

See Also: 2F/42h

**INT 2Fh Function 1229h****DOS 3.3+****READ FROM FILE**

Read data from a previously opened file.

**Call With**

AX 1229h

BX file handle

CX number of bytes to read

DS/DX pointer to buffer

SS DOS DS

**Returns**

as for 2F/3Fh

**Note:** This function can be called to 2F/3Fh, but it is only used when already inside a DOS function call.

This function is used by NLSFUNC to access COUNTRY SYS when invoked by the DOS kernel.

See Also: 2F/3Fh, 2F/1226h

**INT 2Fh Function 122Ah****DOS 3.3+****SET FASTOPEN ENTRY POINT**

Specify the address(es) of the handler(s) for the FASTOPEN requests, each.

## Call With

AX 122Ah  
 BX entry point to set 0001h or 0002h  
 SS CS:0000h or CS:0001h

## Returns

Note: The entry point number is stored under DOS 3.30 entry points are set to 0000h or 0001h by the DOS 4.01 FASTOPEN.

**DOS 3.30+ FASTOPEN is called with**

IF 01h *user*  
 CX *appears to be offset*  
 DI *appears to offset*  
 SI *offset in DOS IPS of filename*  
 IF 02h *unknown*  
 IF 03h *user*  
 SI *offset in DOS IPS of filename*  
 IF 04h *user*  
 AX *subfunction: 00h 01h 02h*  
 FS:DI *points to unknown item*  
 CX *unknown: instructions 01h and 03h only*

## Returns

CF set on error or not successful

Note: The DOS 4.01 FASTOPEN called with additional unknown arguments.

**NT 2Fh Function 122Bh****DOS 3.3+****IOCTL**

Ilocalls an IOCTL control function from within the network redirector.

## Call With

AX Fixed: 122Bh from 21 44h  
 SS DOS IPS  
 add four arguments as appropriate for 21 44h

## Returns

AX 21 44h

Note: The IOCTL control function is called with the same arguments as the DOS 3.30 IOCTL control function.

For net and to fix the frame pointer is restored.

Note: The IOCTL control function is called with the DOS kernel arguments.

**INT 2Fh Function 122Ch****DOS 3.3+****GET DEVICE CHAIN**

IF 01h *user* *offset in DOS IPS of filename*

## Call With

AX CS:0000h

## Returns

AX *offset in DOS IPS of filename* *SI points to device chain*

See Also 2 52h



**INT 2Fh Function 122Dh****DOS 3.3+****GET EXTENDED ERROR CODE**

Return the current extended error code

Call With

AX 122Dh

Returns,

AX current extended error code

See Also: 2F/506, 2F/122Ch

**INT 2Fh Function 122Eh****DOS 4+****GET OR SET ERROR TABLE ADDRESSES**

Specify or determine the locations of various tables used to convert error numbers into error messages

Call With

AX 122Eh

DI subfunction

00h get standard DOS error table (errors 00h-12h/506-516)

Returns,

ES:DI pointer to error table

01h set standard DOS error table

ES:DI pointer to error table

02h get parameter error table (errors 00h/0A6)

Returns

ES:DI pointer to error table

03h set parameter error table

ES:DI pointer to error table

04h get critical SHARE error table (errors 13h/2B6)

Returns

ES:DI pointer to error table

05h set critical SHARE error table

ES:DI pointer to error table

06h get unknown error table

Returns

ES:DI pointer to error table, or 0000h/0000h

07h set unknown error table

ES:DI pointer to error table

08h get error message retriever (see below)

Returns

ES:DI pointer to FAR procedure to fetch error message

09h set unknown error table

ES:DI pointer to error table

**Notes:** If the returned segment on a "get" is 0000h, then the offset specifies the offset of the error message table with a COMMAND.COM, and the procedure returned by DI+08h should be called.

DOS 5.0 COMMAND.COM does not allow setting any of the addresses they are always returned with segment 0000h

See Also: 2F/506, 2F/0500h



It is not necessary to install any of the original BIOS INT 13h BIOS DV key pointer extensions. BIOS 2.00 included special file control disk heads on systems with code type 10000. Call DOS Date Get 10 84 to get the current date and time if the special file control is not installed.

Most DOS system file extensions are selected in BIOS DV through the extensions installed via an INT 13h instruction.

This function presents a unique security loophole for any virus monitoring software which does not trap this call. Some Trojan viruses are known to use it to get the original COM file's pointer.

For detailed discussion of 2F 103 see [Microsoft's Update: DOS Internals](#).

See Also: [INT 13h AH 01h INT 10h](#)

## INT 2Fh Function 1400h

DOS 3.3+

### NLSFUNC.COM - INSTALLATION CHECK

Determine whether NLSFUNC has been loaded.

Called With

AX 1400h

Returns

AX 00h not installed, OK to install  
 01h not installed, not OK  
 11h installed

Notes: 1. For compatibility with some DOS 3.3 versions that support European OS 2.1.3 compatibility, if you are using a system with a localised version of DOS and a localised DOSMIBR, check this function to be certain that NLSFUNC is not installed.

2. Extra information available for MS-DOS 3.31 was not documented in print versions.

## INT 2Fh Function 1400h

European MS-DOS 4.0

### POPUP "CheckPu" INSTALLATION CHECK

Determine whether the OS POPUP dialog box is supported. The POPUP dialog box is supported in MS-DOS 4.00 and later versions of NLSFUNC. Check for MS-DOS 3.31 or later in DOS 3.31 systems installed after European MS-DOS 4.0.

Call With

AX 1400h

Returns

AX 1111h if installed  
 0x00000000 if not installed or successful  
 CF clear if successful  
 CF set on error  
 AX error code

0002h invalid function

0004h an error occurred

Note: 1. Call 1401h to be certain of successful program see 2F 1401h for compatibility details.

See Also: 2F 1401h, 2F 1402h, 2F 1403h

## INT 2Fh Function 1401h

DOS 3.3+

### NLSFUNC.COM CHANGE CODE PAGE

Select a new code page as the default.

Called With

AX 1401h

DSI pointer to internal code page structure - see below  
 BX new code page  
 DX country code

Returns  
 AX status  
   00h successful  
   c.w. DOS error code

Note: This function is called by the DOS 3.3- kernel on 21-38.  
 See Also: 21-38

### Format of DOS 3.30 internal code page structure

Offset	Size	Description
0	WORD	
2	WORD	number of supps
4	5 BYTES	1 1 1 1 1
9	5 BYTES	1 1 1 1 1
14	5 BYTES	1 1 1 1 1
19	5 BYTES	1 1 1 1 1
24	4 BYTES	1 1 1 1

### INT 2Fh Function 1401h

European MS-DOS 4.0

#### POPUP PostPu OPEN/CLOSE POPUP SCREEN

Adjust access to the display from a background program

Call With

AX 1401h  
 DI function: 00h=open, 01h=close  
 DIH save loc:  
   00h block until screen opens  
   01h return error if screen is not available  
   02h urgent - always open screen immediately

Returns

CF clear if successful  
 BX amount of memory needed to save screen and keyboard info  
   0000 - default save location can be used - only if DIH = 02  
 CF set on error

Note: The application is frozen until the popup screen is closed.  
 See Also: 21-1400

### INT 2Fh Function 1402h

DOS 3.3+

#### NLSFUNC COM GET COUNTRY INFO

Called With

AX 1402h  
 BP subfunction - same as M for 21-65h  
 BX code page  
 DX country code  
 DSI pointer to internal code page structure - see 21-1401h  
 ESI pointer to user buffer  
 CX size of user buffer

**Returns:**

AL status  
 00h successful  
 else DOS error code

**Notes:** This call is available only in MS-DOS 3.31 and later versions.

The country code must be supplied by the COUNTRY.MAN installation.

**See Also:** 2f 05h, 2f 1407h, 2f 1404h

**INT 2Fh Function 1402h****European MS-DOS 4.0****POPUP "SavePu" SAVE POPUP SCREEN**

Make a copy of the foreground program's screen with the current buffer and write the copy to a program writes to the screen.

**Call With:**

AX 1402h  
 ES:DI pointer to buffer (0000h to 0000h) to store the copy

**Returns:**

CF clear if successful  
 CF set on error  
 AX error code  
 0001h process does not own screen  
 0004h unknown error  
 0005h invalid pointer

**See Also:** 7f 1400h, 2f 1401h, 2f 1403h

**INT 2Fh Function 1403h****DOS 3.3+****NLSFUNC.COM SET COUNTRY INFO**

Select a new country code as the default.

**Called With:**

AX 1403h  
 DS:SI pointer to buffer (0000h to 0000h) to store the code  
 BX code page  
 DX country code

**Returns:**

AL status

**Note:** This call is available only in MS-DOS 3.31 and later versions.

**See Also:** 2f 06h, 2f 1402h, 2f 1404h

**INT 2Fh Function 1403h****European MS-DOS 4.0****POPUP "RestorePu" RESTORE SCREEN**

Restore the screen to the state it was in before the screen was saved and program write on the screen. The POPUP interface in European MS-DOS 4.0 conflicts with the NLSFUNC interface in MS-DOS 3.31 and later DOS versions available in European MS-DOS 3.0.

**Call With:**

AX 1403h  
 ES:DI pointer to buffer (0000h to 0000h) to store the screen  
 0000h, 0000h for default buffer in POPUP

**Returns:**

CF clear if successful

CF set on error

AX = error code; see AX 1402h "POPUP"

See Also: 2f 1400h 2f 1401c 2f 1402h

## INT 2fh Function 1404h

**DOS 3.3+**

### NLSFUNC.COM GET COUNTRY INFO

Called With

AX = 1404h

BX = code page

DX = country code

DS:SI = pointer to internal code page structure (see 2f 1401h)

ES:DI = pointer to user buffer

Returns

AX = status

Notes: This function is called by the DOS 3.3 keyboard (see 2f 30h)

See Also: 2f 30 2f 1402h 2f 1403h

## INT 2fh Function 14FEh

**DR DOS 5.0**

### NLSFUNC GET EXTENDED COUNTRY INFORMATION

This function returns extended information for a country other than the current country.

Call With

AX = 14FEh

BX = code page 1111h (global code page; see 2f 0602h)

DX = country ID 1111h (current country; see 2f 30h)

ES:DI = pointer to user's buffer of 0100h

CF = 0101h

- 01h get national international info
- 02h get pointer to uppercase table
- 04h get pointer to alternate uppercase table
- 08h get pointer to alternate terminator table
- 0Ch get pointer to collating sequence table
- 0Eh get pointer to Double-Byte Character Set table
- set used to return error if not available

Returns

CF clear if successful

DS:SI = pointer to requested information

CF set if error

Notes: This function is not available in DOS 3.31, DOS 3.3, DOS 3.31a, or DOS 3.31b.

The value of CF is not reset, checked by the DR DOS 5.0 NLSFUNC.

See Also: AX 1411h 2f 065h

### Format of DR DOS COUNTRY SYS file

Offset	Size	Description
00h	01x01h	signature: 0x00000000 / padded with NULs
71h	WORD	signature: FIM In
80h	var	country pointer records







- bit 2: turn 21\_3fh on STDIN into polling loop
- bit 3: trap stack fault in "SYSTEM" to WTN.386
- bit 4: BIOS patch to trap "Insert disk X" to WTN.386

**Returns.**

- AX: B97Ch
- BX: bit mask of patches applied
- DX: A2ABh

0002h remove patches in DOS; ignored by DOS 5.0 kernel

- DX: bit mask of patch requests (see function 0001h)

0003h get size of DOS data structures

- DX: bit mask of request; only one bit can be set
  - bit 0: Current Directory Structure size

**Returns.**

- if supported request
  - AX: B97Ch
  - CX: size in bytes of requested structure
  - DX: A2ABh

etc.

all registers preserved

0004h determine instanced data structures

**Returns.**

- AX: B97Ch if supported
- DX: A2ABh if supported; DOS 5.0 kernel returns 0000h
- BX: bit mask of instanced items
  - bit 0: DS
  - bit 1: SEI
  - bit 2: device list
  - bit 3: DOS swappable data area

0005h get device driver size

FSegment of device driver

**Returns.**

- DX AX: 0000h/0000h sector; not dev. driver segment
- DX AX: A2ABh/B97Ch if successful
- BX CX: size of device driver in bytes

**Notes.** DOSMGR (DOS Manager) will check whether the OEM DOS BIOS data has been instanced via this API and will not perform its own default instancing of the normal DOS EdOS data if so. If this API is not supported, DOSMGR will also try to access instancing data through 21\_1603h. These functions are supported by the DOS 5+ kernel. DOSMGR contains tables of instancing information for earlier versions of DOS.

See Gerald C. Appel's book *DOS Internals* for additional discussions of DOSMGR's behavior and instancing in general. Also see Chapter 1, "GDOSMGR."

See Also: 21\_1603h

**Format of patch table:**

Offset	Size	Description
00h	2 BYTES	DOS version: major:minor
02h	WORD	offset in DOS data segment of "SAVEFS"
04h	WORD	offset in DOS data segment of "SAVEBA"
06h	WORD	offset in DOS data segment of 1st DOS flag
08h	WORD	offset in DOS data segment of 1st ID word
0Ah	WORD	offset in DOS data segment of "ContPatch" table

00000000h      W      00000000h      This is a 32-bit offset in DOS data segment of "UMB HEAD", containing segment of last MCB in conventional memory.

## INT 2Fh Function 160Bh

**Windows 3.1**

### WINDOWS TSR IDENTIFY

Windows 3.1 uses a BIOS interrupt 2Fh, 160Bh. Microsoft does document the function in a Windows 3.1 Source Code Manual Windows Version 3.1 and Digital Language the Microsoft Developer Source Code. See <http://www.microsoft.com/pressrel/pressrel95/0521/0521160B.htm> and <http://www.microsoft.com/pressrel/pressrel95/0521/0521160B.htm>.

## INT 2Fh Function 1684h

**Windows 3.0+**

### GET VXD API ENTRY POINT

Microsoft uses a BIOS interrupt 2Fh, 1684h function in the Windows Device Driver Kit (DDK) to get a pointer to the VxD API entry point of a Windows 3.x VxD. DOS provides a BIOS interrupt 2Fh, 1684h to get a pointer to the API provided by an Enhanced mode virtual device drivers (VxDs).

The Windows 3.1 Source Code Manual Windows Version 3.1 documents the APIs provided by the VxDs. See <http://www.microsoft.com/pressrel/pressrel95/0521/05211684.htm>. Windows programs use the base VxD identifier to get a pointer to the VxD API entry point. The DOS 3.x mode DOS programs use the VxD identifier to get a pointer to the VxD API entry point. The DOS 3.x mode DOS programs use the VxD identifier to get a pointer to the VxD API entry point. By loop through the VxD API entry point, the VxDs provide APIs in which mode. The following table lists the known VxDs.

0003h	VPI.D	Virtual PIO Device	V86, PM
0005h	VT.D	Virtual Timer Device	V86, PM
0009h	Relbnd	Ctrl-Alt-Del virtualization	PM
000Ah	V.D	Virtual Display Device	PM
000Ch	V.M.D	Virtual Mouse Device	V86, PM
000Dh	V.K.D	Virtual Keyboard Device	PM
000Eh	V.C.D	Virtual CMM Device	PM
000Fh	000Fh	DOS 3.x mode DOS extender	V86
0017h	SHR.H	KERNEL-VxD interface	PM
001Ch	LoadH	memory manager	V86
001Dh	WINDBUG	low-level debugging	PM
0021h	Pagefile	demand-paged swap device	PM
0042h	VTDAPI	Multi-media timer	PM

For more information on VxD APIs, see Andrew Tanenbaum's <http://www.winternut.com/~tanenbaum/386/vxd.html> "Virtual Device Drivers for Windows" VxDs. See <http://www.microsoft.com/pressrel/pressrel95/0521/05211684.htm> which also provides a handy list of VxD IDs.

## INT 2Fh Function 168Ah

**DPMI**

### DPMI GET VENDOR SPECIFIC API ENTRY POINT

The DPMI 3.0 specification defines a vendor-specific DPMI, the DOS Protected Mode Interface. Windows 3.1 uses a BIOS interrupt 2Fh, 168Ah to get a pointer to the DPMI service. Microsoft uses a BIOS interrupt 2Fh, 168Ah to get a pointer to the DPMI service. Microsoft uses a BIOS interrupt 2Fh, 168Ah to get a pointer to the DPMI service. For more information, see Matt Pietrek's <http://www.microsoft.com/pressrel/pressrel95/0521/05211684.htm> "Windows Internals, chapter 1".

**INT 2Fh Function 17h****Windows 2.0+****WINDOWS CLIPBOARD API FOR DOS PROGRAMS**

Determine whether Windows 2.0 is installed, and if so, call the Windows API function `WinClipboardOpen` to start using the Windows clipboard. For more information, see the Microsoft KnowledgeBase article "The services are..."

1700h	Identify WinOldAp Version
1701h	Open Clipboard
1702h	Empty Clipboard
1703h	Set Clipboard Data
1704h	Get Clipboard Data Size
1705h	Get Clipboard Data
1708h	Close Clipboard
1709h	Compact Clipboard
170Ah	Get Device Capabilities

For more information, see the article "Clipboard" in *Windows 2.0*.

**INT 2Fh Function 1900h****DOS 4.x only****SHELLB.COM INSTALLATION CHECK**

Determine whether SHELLB has been loaded.

Call With	
AX	1900h
Returns	
AL	00h if not installed 01h if installed

**INT 2Fh Function 1901h****DOS 4.x only****SHELLB.COM SHELLC.EXE INTERFACE**

Inform SHELLB of SHELLC's address, and return the location of a workspace for SHELLC.

Call With	
AX	1901h
BI	00h if SHELLC transient 01h if SHELLC resident
DS:DI	segment address of SHELLC.COM
Returns	
ES	segment address of SHELLC.COM
SI	offset of SHELLC.COM
Note	SHELLC.COM = SHELLC.EXE if not in DOS 4.x file

**INT 2Fh Function 1902h****DOS 4.x only****SHELLB.COM COMMAND.COM INTERFACE**

Call the `COMMAND.COM` interface to execute a batch file from the current batch file.

Call With	
AX	1902h
ES	segment address of SHELLC.COM
DI	filename element (uppercase)
DX	pointer to buffer for results
Returns	
AL	00h failed, other

to the first available byte of the PSN. If no space is available, then the command is not executed. If the command is not executed, then the command is not executed.

If no more Program Start Commands are available:

- AL FFh success, then  
 memory at DS:(DX+1) contains filled as  
 DX+1 BYTEs count of bytes of PNC  
 DX+2 N BYTEs Program Start Command text  
 BYE 00h terminator

**Note:** A VMEXIT to M... can occur at any time, and may occur while a command is in a PNC. If a VMEXIT occurs, the command is not executed. The SHELLB program's Program Start Commands from its workspace.

The SHELLB program's PNCs are not executed. The SHELLB program's PNCs are not executed. The SHELLB program's PNCs are not executed. The SHELLB program's PNCs are not executed.

The SHELLB program's PNCs are not executed. The SHELLB program's PNCs are not executed. The SHELLB program's PNCs are not executed. The SHELLB program's PNCs are not executed.

### INT 2Fh Function 1903h

**DOS 4.x only**

#### SHELLB COM COMMAND.COM INTERFACE

Retrieves the current batch file name and the current batch file.

- Call With**  
 AX 1903h  
 DS:DI pointer to AMIZ batch file name as for 2Fh 1902h

- Returns**  
 AX 1903h if successful, or 00h if SHELLB is a batch file  
 AL 00h if it does not

### INT 2Fh Function 1904h

**DOS 4.x only**

#### SHELLB COM SHELLB TRANSIENT TO TSR Interface

Retrieves the current batch file name and the current batch file.

- Call With**  
 AX 1904h
- Returns**  
 DS:DI pointer to name of current shell batch file  
 WORDS number of bytes of name following  
 BYTEs 8 max uppercase name of shell batch file

### INT 2Fh Function 1A00h

**DOS 4+**

#### ANSI SYS INSTALLATION CHECK

Determine whether ANSI SYS is present.

- Call With**  
 AX 1A00h
- Returns**  
 AL FFh if installed

**Note:** This function is not implemented in DOS 3.11. It was introduced for DOS 4.x.

### INT 2Fh Function 1A01h

**DOS 4+**

#### ANSI SYS GET/SET DISPLAY INFORMATION

It appears to be the ANSI SYS interface to ANSI SYS.

**Call With**

AX 1A01h  
 CL function  
 5H for SI  
 7H for C6

INT 1A01h pointer to parameter block, see 21 1A01h

**Returns**

CF set on error  
 AX error code (many non standard)  
 CF clear if successful  
 AX destroyed

See Also: 40h, 21 1A02h

**INT 2Fh Function 1A02h****DOS 4+****ANSI SYS MISCELLANEOUS REQUESTS**

Get or set miscellaneous ANSI SYS flags

**Call With**

AX 1A02h  
 DS:DX pointer to parameter block, see below

Note: DOS 5+ chains to the previous handler if AX = 02h.

See Also: 21 1A01h

**Format of parameter block**

Offset	Size	Description
00h	BYTE	subfunction 00h: set/reset interlock 01h: get I flag interlock state 00h: reset 01h: set
01h	BYTE	interlock state 00h: reset 01h: set
02h	BYTE	returned 00h if I not in effect 01h if I in effect

**INT 2Fh Function 1B00h****DOS 4+****XMAZEMS SYS INSTALLATION CHECK**

Determine whether XMAZEMS SYS has been loaded

**Call With**

AX 1B00h

**Returns**

AX 1Bh if installed

Note: 1. XMAZEMS SYS is installed if the file XMAZEMS.SYS exists in the root of the drive. This check then hooks onto INT 67. All other checks are made by the user.

See Also: INT 2F A11-110h

**INT 2Fh Function 1Bh****DOS 4+****XMAZEMS SYS GET HIDDEN FRAME INFORMATION**

Determine whether XMAZEMS SYS has been loaded

**Call With**

AH 1Fh  
 AL nonzero  
 DI hidden physical page number

**Returns**

AX EFFFFh if successful  
 0000h if CHK failed  
 ES segment of page 0000h  
 DI physical page number

**Note** INT 2Fh is not available in DOS 3.0. In DOS 3.11, INT 2Fh is available only in real mode.

FASTOPIN makes this call with AL = 1Fh.  
 See Also: 21-1, 100

**INT 2Fh Function 2300h****DR-DOS 5.0****GRAFTABL INSTALLATION CHECK**

Determine whether the DR-DOS GRAFTABL has been installed.

**Call With**

AX 2300h

**Returns**

AH 1Fh

**Note** This is always successful. The usual format of returning AH = 1Fh if installed.  
 See Also: 15-21, AH 23h

**INT 2Fh Function 23h****DR-DOS 5.0****GRAFTABL GET GRAPHICS DATA**

Determine the location of the graphics data used by GRAFTABL.

**Call With**

AH 23h  
 AL nonzero

**Returns**

AX 1Fh  
 ES:BX pointer to graphics data

See Also: 21-2, 200h

**INT 2Fh Function 43h****XMS****XMS 3.0 SERVICES**

Microsoft's XMS 3.0 services are available in the XMS30.SYS file. XMS 3.0 services are not available if you do not have a version 3.0 XMS.SYS file. Microsoft has discretely placed an XMS30.TXT file on the Compaq diskette that comes with XMS 3.0 binaries.

8Bh Query any Free Extended Mem  
 89h Allocate any Extended Memory Block  
 8Fh Get Extended MEMHELDs  
 8Eh Release any Extended Memory

For more information, see Microsoft's XMS30.TXT, or the interrupt list on disk. The XMS 3.0 services are available in the XMS30.SYS file. XMS 3.0 services are not available if you do not have a version 3.0 XMS.SYS file. Microsoft has discretely placed an XMS30.TXT file on the Compaq diskette that comes with XMS 3.0 binaries.

For more information, see Microsoft's XMS30.TXT, or the interrupt list on disk.

---

**INT 2Fh Function 4601h** **DOS 5+**  
**KERNEL UNKNOWN FUNCTION**

The purpose of this function has not been determined.

**Call With**  
 AX = 4601h

**Returns**

AX = 0

**Note** This function copies the MCB following the caller's PSP memory block into the BIOS data word 00.

**See Also:** 4602h

---

**INT 2Fh Function 4602h** **DOS 5+**  
**KERNEL UNKNOWN FUNCTION**

The purpose of this function has not been determined.

**Call With**  
 AX = 4602h

**Returns**

AX = 0

**Note** This function copies the MCB following the caller's BIOS data word 00 into MCB 4 following caller's PSP memory block.

**See Also:** 4601h

---

**INT 2Fh Function 4A00h** **DOS 5+**  
**SINGLE FLOPPY LOGICAL DRIVE CHANGE NOTIFICATION**

This function notifies the user of a floppy disk change on the floppy drive assigned to the only floppy drive controller on the system.

**Called With**  
 AX = 4A00h  
 CX = 0000h  
 DH = new drive number  
 DI = current drive number

**Returns**

CX = FFFh to skip "Insert diskette for drive X" message

**Note** This function is called by MS-DOS 5.03SYSV.SYS when displaying the message "Insert diskette for drive X" (where X is A or B).

---

**INT 2Fh Function 4A01h** **DOS 5+**  
**QUERY FREE HMA SPACE**

This function queries the amount of free space in the Memory Area (MA) M64h-M6Fh.

**Call With**  
 AX = 4A01h

**Returns**

BP = number of bytes of free space in MA (0000h = 0 bytes of free HMA; 0001h = 1 byte of free HMA; 11111111h = 16,777,215 bytes of free MA)

**Note** This function is called by Windows 3.11DOSX.EXE.

**See Also:** 4A02h, 4A03h, 4A04h, 4A05h, 4A06h, 4A07h, 4A08h, 4A09h, 4A0Ah, 4A0Bh, 4A0Ch, 4A0Dh, 4A0Eh, 4A0Fh, 4A10h, 4A11h, 4A12h, 4A13h, 4A14h, 4A15h, 4A16h, 4A17h, 4A18h, 4A19h, 4A1Ah, 4A1Bh, 4A1Ch, 4A1Dh, 4A1Eh, 4A1Fh, 4A20h, 4A21h, 4A22h, 4A23h, 4A24h, 4A25h, 4A26h, 4A27h, 4A28h, 4A29h, 4A2Ah, 4A2Bh, 4A2Ch, 4A2Dh, 4A2Eh, 4A2Fh, 4A30h, 4A31h, 4A32h, 4A33h, 4A34h, 4A35h, 4A36h, 4A37h, 4A38h, 4A39h, 4A3Ah, 4A3Bh, 4A3Ch, 4A3Dh, 4A3Eh, 4A3Fh, 4A40h, 4A41h, 4A42h, 4A43h, 4A44h, 4A45h, 4A46h, 4A47h, 4A48h, 4A49h, 4A4Ah, 4A4Bh, 4A4Ch, 4A4Dh, 4A4Eh, 4A4Fh, 4A50h, 4A51h, 4A52h, 4A53h, 4A54h, 4A55h, 4A56h, 4A57h, 4A58h, 4A59h, 4A5Ah, 4A5Bh, 4A5Ch, 4A5Dh, 4A5Eh, 4A5Fh, 4A60h, 4A61h, 4A62h, 4A63h, 4A64h, 4A65h, 4A66h, 4A67h, 4A68h, 4A69h, 4A6Ah, 4A6Bh, 4A6Ch, 4A6Dh, 4A6Eh, 4A6Fh, 4A70h, 4A71h, 4A72h, 4A73h, 4A74h, 4A75h, 4A76h, 4A77h, 4A78h, 4A79h, 4A7Ah, 4A7Bh, 4A7Ch, 4A7Dh, 4A7Eh, 4A7Fh, 4A80h, 4A81h, 4A82h, 4A83h, 4A84h, 4A85h, 4A86h, 4A87h, 4A88h, 4A89h, 4A8Ah, 4A8Bh, 4A8Ch, 4A8Dh, 4A8Eh, 4A8Fh, 4A90h, 4A91h, 4A92h, 4A93h, 4A94h, 4A95h, 4A96h, 4A97h, 4A98h, 4A99h, 4A9Ah, 4A9Bh, 4A9Ch, 4A9Dh, 4A9Eh, 4A9Fh, 4AA0h, 4AA1h, 4AA2h, 4AA3h, 4AA4h, 4AA5h, 4AA6h, 4AA7h, 4AA8h, 4AA9h, 4AAAh, 4AA Bh, 4AACh, 4AA Dh, 4AA Eh, 4AA Fh, 4AB0h, 4AB1h, 4AB2h, 4AB3h, 4AB4h, 4AB5h, 4AB6h, 4AB7h, 4AB8h, 4AB9h, 4ABAh, 4AB Bh, 4ABCh, 4AB Dh, 4AB Eh, 4AB Fh, 4AC0h, 4AC1h, 4AC2h, 4AC3h, 4AC4h, 4AC5h, 4AC6h, 4AC7h, 4AC8h, 4AC9h, 4ACAh, 4AC Bh, 4ACCh, 4AC Dh, 4AC Eh, 4AC Fh, 4AD0h, 4AD1h, 4AD2h, 4AD3h, 4AD4h, 4AD5h, 4AD6h, 4AD7h, 4AD8h, 4AD9h, 4ADAh, 4AD Bh, 4ADCh, 4AD Dh, 4AD Eh, 4AD Fh, 4AE0h, 4AE1h, 4AE2h, 4AE3h, 4AE4h, 4AE5h, 4AE6h, 4AE7h, 4AE8h, 4AE9h, 4AEAh, 4AE Bh, 4AEC h, 4AEDh, 4AEEh, 4AEFh, 4AF0h, 4AF1h, 4AF2h, 4AF3h, 4AF4h, 4AF5h, 4AF6h, 4AF7h, 4AF8h, 4AF9h, 4AFAh, 4AF Bh, 4AFCh, 4AF Dh, 4AF Eh, 4AF Fh, 4B00h, 4B01h, 4B02h, 4B03h, 4B04h, 4B05h, 4B06h, 4B07h, 4B08h, 4B09h, 4B0Ah, 4B0 Bh, 4B0Ch, 4B0 Dh, 4B0 Eh, 4B0 Fh, 4B10h, 4B11h, 4B12h, 4B13h, 4B14h, 4B15h, 4B16h, 4B17h, 4B18h, 4B19h, 4B1Ah, 4B1 Bh, 4B1Ch, 4B1 Dh, 4B1 Eh, 4B1 Fh, 4B20h, 4B21h, 4B22h, 4B23h, 4B24h, 4B25h, 4B26h, 4B27h, 4B28h, 4B29h, 4B2Ah, 4B2 Bh, 4B2Ch, 4B2 Dh, 4B2 Eh, 4B2 Fh, 4B30h, 4B31h, 4B32h, 4B33h, 4B34h, 4B35h, 4B36h, 4B37h, 4B38h, 4B39h, 4B3Ah, 4B3 Bh, 4B3Ch, 4B3 Dh, 4B3 Eh, 4B3 Fh, 4B40h, 4B41h, 4B42h, 4B43h, 4B44h, 4B45h, 4B46h, 4B47h, 4B48h, 4B49h, 4B4Ah, 4B4 Bh, 4B4Ch, 4B4 Dh, 4B4 Eh, 4B4 Fh, 4B50h, 4B51h, 4B52h, 4B53h, 4B54h, 4B55h, 4B56h, 4B57h, 4B58h, 4B59h, 4B5Ah, 4B5 Bh, 4B5Ch, 4B5 Dh, 4B5 Eh, 4B5 Fh, 4B60h, 4B61h, 4B62h, 4B63h, 4B64h, 4B65h, 4B66h, 4B67h, 4B68h, 4B69h, 4B6Ah, 4B6 Bh, 4B6Ch, 4B6 Dh, 4B6 Eh, 4B6 Fh, 4B70h, 4B71h, 4B72h, 4B73h, 4B74h, 4B75h, 4B76h, 4B77h, 4B78h, 4B79h, 4B7Ah, 4B7 Bh, 4B7Ch, 4B7 Dh, 4B7 Eh, 4B7 Fh, 4B80h, 4B81h, 4B82h, 4B83h, 4B84h, 4B85h, 4B86h, 4B87h, 4B88h, 4B89h, 4B8Ah, 4B8 Bh, 4B8Ch, 4B8 Dh, 4B8 Eh, 4B8 Fh, 4B90h, 4B91h, 4B92h, 4B93h, 4B94h, 4B95h, 4B96h, 4B97h, 4B98h, 4B99h, 4B9Ah, 4B9 Bh, 4B9Ch, 4B9 Dh, 4B9 Eh, 4B9 Fh, 4BA0h, 4BA1h, 4BA2h, 4BA3h, 4BA4h, 4BA5h, 4BA6h, 4BA7h, 4BA8h, 4BA9h, 4BAAh, 4BABh, 4BAC h, 4BADh, 4BAEh, 4BAFh, 4BB0h, 4BB1h, 4BB2h, 4BB3h, 4BB4h, 4BB5h, 4BB6h, 4BB7h, 4BB8h, 4BB9h, 4BBAh, 4BB Bh, 4BBCh, 4BB Dh, 4BB Eh, 4BB Fh, 4BC0h, 4BC1h, 4BC2h, 4BC3h, 4BC4h, 4BC5h, 4BC6h, 4BC7h, 4BC8h, 4BC9h, 4BCAh, 4BC Bh, 4BCCh, 4BC Dh, 4BC Eh, 4BC Fh, 4BD0h, 4BD1h, 4BD2h, 4BD3h, 4BD4h, 4BD5h, 4BD6h, 4BD7h, 4BD8h, 4BD9h, 4BDAh, 4BD Bh, 4BDCh, 4BD Dh, 4BD Eh, 4BD Fh, 4BE0h, 4BE1h, 4BE2h, 4BE3h, 4BE4h, 4BE5h, 4BE6h, 4BE7h, 4BE8h, 4BE9h, 4BEAh, 4BE Bh, 4BECh, 4BE Dh, 4BE Eh, 4BE Fh, 4BF0h, 4BF1h, 4BF2h, 4BF3h, 4BF4h, 4BF5h, 4BF6h, 4BF7h, 4BF8h, 4BF9h, 4BFAh, 4BF Bh, 4BFCh, 4BF Dh, 4BF Eh, 4BF Fh, 4C00h, 4C01h, 4C02h, 4C03h, 4C04h, 4C05h, 4C06h, 4C07h, 4C08h, 4C09h, 4C0Ah, 4C0 Bh, 4C0Ch, 4C0 Dh, 4C0 Eh, 4C0 Fh, 4C10h, 4C11h, 4C12h, 4C13h, 4C14h, 4C15h, 4C16h, 4C17h, 4C18h, 4C19h, 4C1Ah, 4C1 Bh, 4C1Ch, 4C1 Dh, 4C1 Eh, 4C1 Fh, 4C20h, 4C21h, 4C22h, 4C23h, 4C24h, 4C25h, 4C26h, 4C27h, 4C28h, 4C29h, 4C2Ah, 4C2 Bh, 4C2Ch, 4C2 Dh, 4C2 Eh, 4C2 Fh, 4C30h, 4C31h, 4C32h, 4C33h, 4C34h, 4C35h, 4C36h, 4C37h, 4C38h, 4C39h, 4C3Ah, 4C3 Bh, 4C3Ch, 4C3 Dh, 4C3 Eh, 4C3 Fh, 4C40h, 4C41h, 4C42h, 4C43h, 4C44h, 4C45h, 4C46h, 4C47h, 4C48h, 4C49h, 4C4Ah, 4C4 Bh, 4C4Ch, 4C4 Dh, 4C4 Eh, 4C4 Fh, 4C50h, 4C51h, 4C52h, 4C53h, 4C54h, 4C55h, 4C56h, 4C57h, 4C58h, 4C59h, 4C5Ah, 4C5 Bh, 4C5Ch, 4C5 Dh, 4C5 Eh, 4C5 Fh, 4C60h, 4C61h, 4C62h, 4C63h, 4C64h, 4C65h, 4C66h, 4C67h, 4C68h, 4C69h, 4C6Ah, 4C6 Bh, 4C6Ch, 4C6 Dh, 4C6 Eh, 4C6 Fh, 4C70h, 4C71h, 4C72h, 4C73h, 4C74h, 4C75h, 4C76h, 4C77h, 4C78h, 4C79h, 4C7Ah, 4C7 Bh, 4C7Ch, 4C7 Dh, 4C7 Eh, 4C7 Fh, 4C80h, 4C81h, 4C82h, 4C83h, 4C84h, 4C85h, 4C86h, 4C87h, 4C88h, 4C89h, 4C8Ah, 4C8 Bh, 4C8Ch, 4C8 Dh, 4C8 Eh, 4C8 Fh, 4C90h, 4C91h, 4C92h, 4C93h, 4C94h, 4C95h, 4C96h, 4C97h, 4C98h, 4C99h, 4C9Ah, 4C9 Bh, 4C9Ch, 4C9 Dh, 4C9 Eh, 4C9 Fh, 4CA0h, 4CA1h, 4CA2h, 4CA3h, 4CA4h, 4CA5h, 4CA6h, 4CA7h, 4CA8h, 4CA9h, 4CAAh, 4CABh, 4CAC h, 4CADh, 4CAEh, 4CAFh, 4CB0h, 4CB1h, 4CB2h, 4CB3h, 4CB4h, 4CB5h, 4CB6h, 4CB7h, 4CB8h, 4CB9h, 4CBAh, 4CB Bh, 4CBCh, 4CB Dh, 4CB Eh, 4CB Fh, 4CC0h, 4CC1h, 4CC2h, 4CC3h, 4CC4h, 4CC5h, 4CC6h, 4CC7h, 4CC8h, 4CC9h, 4CCAh, 4CC Bh, 4CCCh, 4CC Dh, 4CC Eh, 4CC Fh, 4CD0h, 4CD1h, 4CD2h, 4CD3h, 4CD4h, 4CD5h, 4CD6h, 4CD7h, 4CD8h, 4CD9h, 4CDAh, 4CD Bh, 4CDCh, 4CD Dh, 4CD Eh, 4CD Fh, 4CE0h, 4CE1h, 4CE2h, 4CE3h, 4CE4h, 4CE5h, 4CE6h, 4CE7h, 4CE8h, 4CE9h, 4CEAh, 4CE Bh, 4CECh, 4CE Dh, 4CE Eh, 4CE Fh, 4CF0h, 4CF1h, 4CF2h, 4CF3h, 4CF4h, 4CF5h, 4CF6h, 4CF7h, 4CF8h, 4CF9h, 4CFAh, 4CF Bh, 4CFCh, 4CF Dh, 4CF Eh, 4CF Fh, 4D00h, 4D01h, 4D02h, 4D03h, 4D04h, 4D05h, 4D06h, 4D07h, 4D08h, 4D09h, 4D0Ah, 4D0 Bh, 4D0Ch, 4D0 Dh, 4D0 Eh, 4D0 Fh, 4D10h, 4D11h, 4D12h, 4D13h, 4D14h, 4D15h, 4D16h, 4D17h, 4D18h, 4D19h, 4D1Ah, 4D1 Bh, 4D1Ch, 4D1 Dh, 4D1 Eh, 4D1 Fh, 4D20h, 4D21h, 4D22h, 4D23h, 4D24h, 4D25h, 4D26h, 4D27h, 4D28h, 4D29h, 4D2Ah, 4D2 Bh, 4D2Ch, 4D2 Dh, 4D2 Eh, 4D2 Fh, 4D30h, 4D31h, 4D32h, 4D33h, 4D34h, 4D35h, 4D36h, 4D37h, 4D38h, 4D39h, 4D3Ah, 4D3 Bh, 4D3Ch, 4D3 Dh, 4D3 Eh, 4D3 Fh, 4D40h, 4D41h, 4D42h, 4D43h, 4D44h, 4D45h, 4D46h, 4D47h, 4D48h, 4D49h, 4D4Ah, 4D4 Bh, 4D4Ch, 4D4 Dh, 4D4 Eh, 4D4 Fh, 4D50h, 4D51h, 4D52h, 4D53h, 4D54h, 4D55h, 4D56h, 4D57h, 4D58h, 4D59h, 4D5Ah, 4D5 Bh, 4D5Ch, 4D5 Dh, 4D5 Eh, 4D5 Fh, 4D60h, 4D61h, 4D62h, 4D63h, 4D64h, 4D65h, 4D66h, 4D67h, 4D68h, 4D69h, 4D6Ah, 4D6 Bh, 4D6Ch, 4D6 Dh, 4D6 Eh, 4D6 Fh, 4D70h, 4D71h, 4D72h, 4D73h, 4D74h, 4D75h, 4D76h, 4D77h, 4D78h, 4D79h, 4D7Ah, 4D7 Bh, 4D7Ch, 4D7 Dh, 4D7 Eh, 4D7 Fh, 4D80h, 4D81h, 4D82h, 4D83h, 4D84h, 4D85h, 4D86h, 4D87h, 4D88h, 4D89h, 4D8Ah, 4D8 Bh, 4D8Ch, 4D8 Dh, 4D8 Eh, 4D8 Fh, 4D90h, 4D91h, 4D92h, 4D93h, 4D94h, 4D95h, 4D96h, 4D97h, 4D98h, 4D99h, 4D9Ah, 4D9 Bh, 4D9Ch, 4D9 Dh, 4D9 Eh, 4D9 Fh, 4DA0h, 4DA1h, 4DA2h, 4DA3h, 4DA4h, 4DA5h, 4DA6h, 4DA7h, 4DA8h, 4DA9h, 4DAAh, 4DABh, 4DAC h, 4DADh, 4DAEh, 4DAFh, 4DB0h, 4DB1h, 4DB2h, 4DB3h, 4DB4h, 4DB5h, 4DB6h, 4DB7h, 4DB8h, 4DB9h, 4DBAh, 4DB Bh, 4DBCh, 4DB Dh, 4DB Eh, 4DB Fh, 4DC0h, 4DC1h, 4DC2h, 4DC3h, 4DC4h, 4DC5h, 4DC6h, 4DC7h, 4DC8h, 4DC9h, 4DCAh, 4DC Bh, 4DCCh, 4DC Dh, 4DC Eh, 4DC Fh, 4DD0h, 4DD1h, 4DD2h, 4DD3h, 4DD4h, 4DD5h, 4DD6h, 4DD7h, 4DD8h, 4DD9h, 4DDAh, 4DD Bh, 4DDCh, 4DD Dh, 4DD Eh, 4DD Fh, 4DE0h, 4DE1h, 4DE2h, 4DE3h, 4DE4h, 4DE5h, 4DE6h, 4DE7h, 4DE8h, 4DE9h, 4DEAh, 4DE Bh, 4DECh, 4DE Dh, 4DE Eh, 4DE Fh, 4DF0h, 4DF1h, 4DF2h, 4DF3h, 4DF4h, 4DF5h, 4DF6h, 4DF7h, 4DF8h, 4DF9h, 4DFAh, 4DF Bh, 4DFCh, 4DF Dh, 4DF Eh, 4DF Fh, 4E00h, 4E01h, 4E02h, 4E03h, 4E04h, 4E05h, 4E06h, 4E07h, 4E08h, 4E09h, 4E0Ah, 4E0 Bh, 4E0Ch, 4E0 Dh, 4E0 Eh, 4E0 Fh, 4E10h, 4E11h, 4E12h, 4E13h, 4E14h, 4E15h, 4E16h, 4E17h, 4E18h, 4E19h, 4E1Ah, 4E1 Bh, 4E1Ch, 4E1 Dh, 4E1 Eh, 4E1 Fh, 4E20h, 4E21h, 4E22h, 4E23h, 4E24h, 4E25h, 4E26h, 4E27h, 4E28h, 4E29h, 4E2Ah, 4E2 Bh, 4E2Ch, 4E2 Dh, 4E2 Eh, 4E2 Fh, 4E30h, 4E31h, 4E32h, 4E33h, 4E34h, 4E35h, 4E36h, 4E37h, 4E38h, 4E39h, 4E3Ah, 4E3 Bh, 4E3Ch, 4E3 Dh, 4E3 Eh, 4E3 Fh, 4E40h, 4E41h, 4E42h, 4E43h, 4E44h, 4E45h, 4E46h, 4E47h, 4E48h, 4E49h, 4E4Ah, 4E4 Bh, 4E4Ch, 4E4 Dh, 4E4 Eh, 4E4 Fh, 4E50h, 4E51h, 4E52h, 4E53h, 4E54h, 4E55h, 4E56h, 4E57h, 4E58h, 4E59h, 4E5Ah, 4E5 Bh, 4E5Ch, 4E5 Dh, 4E5 Eh, 4E5 Fh, 4E60h, 4E61h, 4E62h, 4E63h, 4E64h, 4E65h, 4E66h, 4E67h, 4E68h, 4E69h, 4E6Ah, 4E6 Bh, 4E6Ch, 4E6 Dh, 4E6 Eh, 4E6 Fh, 4E70h, 4E71h, 4E72h, 4E73h, 4E74h, 4E75h, 4E76h, 4E77h, 4E78h, 4E79h, 4E7Ah, 4E7 Bh, 4E7Ch, 4E7 Dh, 4E7 Eh, 4E7 Fh, 4E80h, 4E81h, 4E82h, 4E83h, 4E84h, 4E85h, 4E86h, 4E87h, 4E88h, 4E89h, 4E8Ah, 4E8 Bh, 4E8Ch, 4E8 Dh, 4E8 Eh, 4E8 Fh, 4E90h, 4E91h, 4E92h, 4E93h, 4E94h, 4E95h, 4E96h, 4E97h, 4E98h, 4E99h, 4E9Ah, 4E9 Bh, 4E9Ch, 4E9 Dh, 4E9 Eh, 4E9 Fh, 4EA0h, 4EA1h, 4EA2h, 4EA3h, 4EA4h, 4EA5h, 4EA6h, 4EA7h, 4EA8h, 4EA9h, 4EAAh, 4EABh, 4EAC h, 4EADh, 4EA Eh, 4EAFh, 4EB0h, 4EB1h, 4EB2h, 4EB3h, 4EB4h, 4EB5h, 4EB6h, 4EB7h, 4EB8h, 4EB9h, 4EBAh, 4EB Bh, 4EBCh, 4EB Dh, 4EB Eh, 4EB Fh, 4EC0h, 4EC1h, 4EC2h, 4EC3h, 4EC4h, 4EC5h, 4EC6h, 4EC7h, 4EC8h, 4EC9h, 4ECAh, 4EC Bh, 4ECCh, 4EC Dh, 4EC Eh, 4EC Fh, 4ED0h, 4ED1h, 4ED2h, 4ED3h, 4ED4h, 4ED5h, 4ED6h, 4ED7h, 4ED8h, 4ED9h, 4EDAh, 4ED Bh, 4EDCh, 4ED Dh, 4ED Eh, 4ED Fh, 4EE0h, 4EE1h, 4EE2h, 4EE3h, 4EE4h, 4EE5h, 4EE6h, 4EE7h, 4EE8h, 4EE9h, 4EEAh, 4EE Bh, 4EECh, 4EE Dh, 4EE Eh, 4EE Fh, 4EF0h, 4EF1h, 4EF2h, 4EF3h, 4EF4h, 4EF5h, 4EF6h, 4EF7h, 4EF8h, 4EF9h, 4EFAh, 4EF Bh, 4EFCh, 4EF Dh, 4EF Eh, 4EF Fh, 4F00h, 4F01h, 4F02h, 4F03h, 4F04h, 4F05h, 4F06h, 4F07h, 4F08h, 4F09h, 4F0Ah, 4F0 Bh, 4F0Ch, 4F0 Dh, 4F0 Eh, 4F0 Fh, 4F10h, 4F11h, 4F12h, 4F13h, 4F14h, 4F15h, 4F16h, 4F17h, 4F18h, 4F19h, 4F1Ah, 4F1 Bh, 4F1Ch, 4F1 Dh, 4F1 Eh, 4F1 Fh, 4F20h, 4F21h, 4F22h, 4F23h, 4F24h, 4F25h, 4F26h, 4F27h, 4F28h, 4F29h, 4F2Ah, 4F2 Bh, 4F2Ch, 4F2 Dh, 4F2 Eh, 4F2 Fh, 4F30h, 4F31h, 4F32h, 4F33h, 4F34h, 4F35h, 4F36h, 4F37h, 4F38h, 4F39h, 4F3Ah, 4F3 Bh, 4F3Ch, 4F3 Dh, 4F3 Eh, 4F3 Fh, 4F40h, 4F41h, 4F42h, 4F43h, 4F44h, 4F45h, 4F46h, 4F47h, 4F48h, 4F49h, 4F4Ah, 4F4 Bh, 4F4Ch, 4F4 Dh, 4F4 Eh, 4F4 Fh, 4F50h, 4F51h, 4F52h, 4F53h, 4F54h, 4F55h, 4F56h, 4F57h, 4F58h, 4F59h, 4F5Ah, 4F5 Bh, 4F5Ch, 4F5 Dh, 4F5 Eh, 4F5 Fh, 4F60h, 4F61h, 4F62h, 4F63h, 4F64h, 4F65h, 4F66h, 4F67h, 4F68h, 4F69h, 4F6Ah, 4F6 Bh, 4F6Ch, 4F6 Dh, 4F6 Eh, 4F6 Fh, 4F70h, 4F71h, 4F72h, 4F73h, 4F74h, 4F75h, 4F76h, 4F77h, 4F78h, 4F79h, 4F7Ah, 4F7 Bh, 4F7Ch, 4F7 Dh, 4F7 Eh, 4F7 Fh, 4F80h, 4F81h, 4F82h, 4F83h, 4F84h, 4F85h, 4F86h, 4F87h, 4F88h, 4F89h, 4F8Ah, 4F8 Bh, 4F8Ch, 4F8 Dh, 4F8 Eh, 4F8 Fh, 4F90h, 4F91h, 4F92h, 4F93h, 4F94h, 4F95h, 4F96h, 4F97h, 4F98h, 4F99h, 4F9Ah, 4F9 Bh, 4F9Ch, 4F9 Dh, 4F9 Eh, 4F9 Fh, 4FA0h, 4FA1h, 4FA2h, 4FA3h, 4FA4h, 4FA5h, 4FA6h, 4FA7h, 4FA8h, 4FA9h, 4FAAh, 4FABh, 4FAC h, 4FADh, 4FAEh, 4FAFh, 4FB0h, 4FB1h, 4FB2h, 4FB3h, 4FB4h, 4FB5h, 4FB6h, 4FB7h, 4FB8h, 4FB9h, 4FBAh, 4FB Bh, 4FBCh, 4FB Dh, 4FB Eh, 4FB Fh, 4FC0h, 4FC1h, 4FC2h, 4FC3h, 4FC4h, 4FC5h, 4FC6h, 4FC7h, 4FC8h, 4FC9h, 4FCAh, 4FC Bh, 4FCCh, 4FC Dh, 4FC Eh, 4FC Fh, 4FD0h, 4FD1h, 4FD2h, 4FD3h, 4FD4h, 4FD5h, 4FD6h, 4FD7h, 4FD8h, 4FD9h, 4FDAh, 4FD Bh, 4FDCh, 4FD Dh, 4FD Eh, 4FD Fh, 4FE0h, 4FE1h, 4FE2h, 4FE3h, 4FE4h, 4FE5h, 4FE6h, 4FE7h, 4FE8h, 4FE9h, 4FEAh, 4FE Bh, 4FECh, 4FE Dh, 4FE Eh, 4FE Fh, 4FF0h, 4FF1h, 4FF2h, 4FF3h, 4FF4h, 4FF5h, 4FF6h, 4FF7h, 4FF8h, 4FF9h, 4FFAh, 4FF Bh, 4FFCh, 4FF Dh, 4FF Eh, 4FF Fh

**See Also:** 4A07h

---

**INT 2Fh Function 4A02h**  
**ALLOCATE HMA SPACE**
**DOS 5.0**

Reserve a portion of the High Memory Area for storing a portion of the TSR's code or data.

Called With

AX = 4A02h  
 BX = number of bytes

Returns

ES:DI = pointer to start of allocated HMA block or FFFFFFFFh

Changes

**NOTE:** If a memory DOS is loaded in the HMA, DOS HIGH in CONSRG.SYS (the object or module called by Word 5.31 DOSX.EXE) must be set to allocate HMA space.

See Also: 2Fh 4A00h

---

**INT 2Fh Function 4A05h**  
**DOSSHELL TASK SWITCHING API**
**DOS 5+**

Called With

AX = 4A05h  
 SI = pointer to command line  
 DI = pointer to shell function

See Also

See Also: 1Fh 4B00h

---

**INT 2Fh function 4A06h**  
**ADJUST MEMORY SIZE**
**DOS 5+**

Called With

AX = 4A06h  
 DX = segment following last byte of conventional mem

Returns

DX = segment following last byte of memory available for use by DOS

**NOTE:** This function is only available in DOS 5.0 and later. For more information, see the *Microsoft Windows 3.11 Programmer's Reference*.

INT 12h See Chapter 3, *DOS Internals*, Chapter 3.

See Also: 1Fh 4A00h

---

**INT 2Fh Function 4A10h Subfn 0000h**  
**Smartdrv Installation Check And Hit Ratios**
**SMARTDRV v4+**

Smartdrv is a disk cache that improves disk access times and is highly effective in reducing the number of disk accesses. Smartdrv is installed in the first free slot listing on disk under "31-4403".



**Call With**

AX 4A10h  
 BX 0000h  
 CX EFBABh v4.1+ see Note

**Returns**

AX BABFh if installed  
 DX:BX cache hits  
 DI:SI cache misses  
 BP version (400-0400h)  
 CX *unknown*

**Notes** Also see SMARTDRV.ALI, included if the cache supports IDE/PC Cache v8.0  
 If DISKSPACE.FLN is installed by SMARTDRV, the user checks its availability by CX:EBXh  
 or by DISKSPACE.FLN. If yes, a error message "Cache at SMARTDRV v4.0 with  
 DoubleSpace" and aborts the caller with 21-4C00h

**See Also** AX 4A10h BX 0001h AX 4A10h BX 0003h AX 4A10h BX 0004h  
 AX 4A10h BX 0005h AX 4A10h BX 0007h AX 4A10h BX 1334h 21-4102h  
 "SMARTDRV" 21-4403h "SMARTDRV"

### INT 2Fh Function 4A10h Subfn 0001h RESET CACHE

**SMARTDRV v4+****Call With**

AX 4A10h  
 BX 0001h

**Returns**

registers unchanged

**Note** This function is also supported by PC Cache v8.0

**See Also** AX 4A10h BX 0000h AX 4A10h BX 0002h

### INT 2Fh Function 4A10h Subfn 0002h FLUSH BUFFERS

**SMARTDRV v4+**

Force all modified data to be written to disk immediately

**Call With**

AX 4A10h  
 BX 0002h

**Returns**

registers unchanged

**Note** This function is also supported by PC Cache v8.0

**See Also** AX 4A10h BX 0000h AX 4A10h BX 0001h

### INT 2Fh Function 4A10h Subfn 0003h STATUS

**SMARTDRV v4+**

Determine the caching status of a specified drive

**Call With**

AX 0000h  
 BX 0002h  
 BP drive # (D-A, I-B, etc.)  
 DI subfunction:  
 00h turn on read cache  
 01h turn on read cache

00000000 read cache  
 00000000 write cache  
 00000000 write cache

**Returns**

XX SMARTDRV.CAB  
 DX File path does not exist

cached  
 size of cache

IO drive # 0-3 A-F

**See Also**

SMARTDRV.CAB

This function is also supported by PC Cache v8.0

See Also XX 4A10h IX 0004h

**INT 2Fh Function 4A10h Subfn 0004h****SMARTDRV v4+****GET CACHE SIZE**

DOS 3.21+

**Call With**

XX 4A10h  
 IX 0004h

**Returns**

XX cache size in bytes  
 IX largest number of items  
 CX size of elements in IX  
 DX number of elements under Windows

**Note:** This function is also supported by PC Cache v8.0

See Also XX 4A10h IX 0000h XX 4A10h IX 0004h

**INT 2Fh Function 4A10h Subfn 0005h****SMARTDRV v4+****GET DOUBLE BUFFER STATUS**

Determine whether or the double buffering code was on.

**Call With**

XX 4A10h  
 IX 0005h  
 BP drive # 0-3 A-F

**Returns**

XX SMARTDRV.CAB  
 IX SMARTDRV.CAB

**See Also**

SMARTDRV.CAB  
 SMARTDRV.CAB

**INT 2Fh Function 4A10h Subfn 0006h****SMARTDRV v4+****CHECK IF DRIVE CACHEABLE**

SMARTDRV v4.15 and later: a startup to determine whether it should cache a particular drive

**Called with:**

XX 4A10h  
 IX 0006h  
 CX drive number 01h - A

**Returns:**

AX = 0006h if drive should not be cached by SMARTDRV

See Also: AX = 4A10h BX = 0000h

**INT 2Fh Function 4A10h Subfn 0007h****SMARTDRV v4+****GET DEVICE DRIVER FOR DRIVE**

The purpose of this function has not yet been determined.

**Call With:**

AX = 4A10h

BX = 0007h

BP = drive number

**Returns:**

DI = *unknown*

FSI = pointer to device driver header for drive

Note: This function is also supported by PC Cache v8.0

See Also: AX = 4A10h BX = 0000h

**INT 2Fh Function 4A10h Subfn 000Ah****SMARTDRV v4+****GET UNKNOWN TABLE POINTER**

The purpose of this function has not yet been determined.

**Call With:**

AX = 4A10h

BX = 000Ah

**Returns:**

ES:AX = pointer to unknown table containing information

ES:BX = pointer to unknown table containing info (see below)

Note: This function is also supported by PC Cache v8.0

See Also: AX = 4A10h BX = 0000h

**Format of data table**

Offset	Size	Description
00h	8 BYTES	<i>unknown</i>
08h	32 BYTES	<i>unknown</i>

**INT 2Fh Function 4A10h Subfn 1234h****SMARTDRV v4+****SIGNAL SERIOUS ERROR**

The purpose of this function has not yet been determined.

**Call With:**

AX = 4A10h

BX = 1234h

Note: This function is also supported by PC Cache v8.0

See Also: AX = 4A10h BX = 0000h

**INT 2Fh Function 4A11h Subfunc FFFeh****DBLSPACE BIN****RELOCATE**

Move the DBLSPACE.BIN driver to its final location in memory.

**Call With:**

AX = 4A11h

BX = FFFeh

ES = segment to which to relocate DBLSPACE.BIN

## Returns

AX = 0000h if successful  
 Other DoubleSpace API calls include functions 3 and 4, which at one point were "reserved" in *MS-DOS 3.11* and documented in *Microsoft Windows Programmers Reference*.  
 See Also: AX 4410h BX 4412h

**INT 2Fh Function 4A11h Subfunc FFFFh  
 GET RELOCATION SIZE**
**DBLSPACE.BIN**

This function of the DBLSPACE.BIN driver requires  
 Call With  
 AX 4A11h  
 BX 4411h

## Returns

AX = number of paragraphs needed by DBLSPACE.BIN  
 See Also: AX 4410h BX 4412h

**INT 2Fh Function 4A13h  
 GET UNKNOWN ENTRY POINTS**
**DBLSPACE.BIN**

The purpose of this call has not been determined.  
 Call With  
 AX 4A13h

## Returns

AX 4440h supported  
 EBX = pointer to entry point record, see below  
 See Also: AX 4410h BX 4412h

**Format of entry point record**

Offset	Size	Description
00h	DWORD	pointer to unknown FAR function
4	DWORD	FAR JUMP instruction to unknown function

**INT 2Fh Function 5500h  
 COMMAND.COM INTERFACE**
**DOS 5+**

This function of the COMMAND.COM interface requires  
 Call With  
 AX 5500h

## Returns

AX 0000h  
 BX = pointer to entry point table  
 See Also: COMMAND.COM

Note: The entry point table in COMMAND.COM is located at the resident portion of the stack and are thus not

**INT 2Fh Function 5600h****DOS 6+****INTERLNK Install Check**

Determine whether Microsoft's device driver for allowing one computer to access another is installed.

**INT 2Fh Function 9400h****Workgroup Connection****MICRO.EXE INSTALLATION CHECK**

Determine whether the MICRO.COM device driver for Workgroup connections is installed.

Call With

AX 9400h

Returns

AX 07h or 08h if installed

Note: 2F 9401-9404 are also used, but their meaning is unknown.

See Also: 0000, AX 9402, AX 9403, AX 9404, 2F 9405, 4C00, 4C01, 4C02, 4C03, 4C04, 4C05

**INT 2Fh Function AC00h****DOS 4.01+****GRAPHICS COM INSTALLATION CHECK**

Determine whether GRAPHICS has been loaded.

Call With

AX AC00h

Returns

AX FEFFh

Note: This function is used by the `GRAPHICS.COM` driver to determine if it is already loaded.

Note: This function is used by the `GRAPHICS.COM` driver to determine if it is already loaded.

that occurred in DOS 4.00.

See Also: 2F 1500h

**INT 2Fh Function AD00h****DOS 3.3+****DISPLAY.SYS INSTALLATION CHECK**

Determine whether DISPLAY.SYS is present.

Call With

AX AD00h

Returns

AX FFh if installed

Note: This function is used by the `DISPLAY.SYS` driver to determine if it is already loaded.

Note: This function is used by the `DISPLAY.SYS` driver to determine if it is already loaded.

**INT 2Fh Function AD01h****DOS 3.3+****DISPLAY.SYS INTERNAL SET ACTIVE CODE PAGE**

Specify which code page DISPLAY.SYS should use.

Call With

AX AD01h

BX new code page

Returns

CF clear if successful

AX 00011

CF set if error (unsupported code page)

AX 00000

See Also: 21\_A302

### INT 2Fh Function AD02h

**DOS 3.3+**

#### DISPLAY SYS INTERNAL GET ACTIVE CODE PAGE

page DISPLAY SYS is currently using

Call With

AX A302h

Returns

BX 1111h (unknown first hardware code page)

CF 0 if successful

BX current code page

See Also: 21\_A300h 21\_A303h

### INT 2Fh Function AD03h

**DOS 3.3+**

#### DISPLAY SYS INTERNAL GET CODE PAGE INFORMATION

Get code page information: the number of code pages supported and the currently available code pages

Call With

AX A303h

ES:DI pointer to buffer for code page information (see below)

CX size of output buffers

Returns

CF set if error (no output)

CF 0 if successful

ES:DI buffer index

See Also: 21\_A300h 21\_A301h

#### Format of DOS 5.0 code page information

Offset	Size	Description
00h	WORD	number of software code pages
02h	WORD	<i>unknown offset</i>
04h	WORD	number of hardware code pages
06h	NWORD	hardware code page numbers

0 1 2 3 4 5 6 7 8 9

### INT 2Fh Function AD10h

**DOS 4.x only**

#### DISPLAY SYS INTERNAL INSTALLATION CHECK

see appendix to DOS 10.x

see function for APPEND then it also returns the DISPLAY SYS version

Call With

AX A310h

*additional arguments if any*

Returns

AX 1111h

BX unknown 0100h, 1000h, 40

### INT 2Fh Function AD10h DISPLAY.SYS INTERNAL - UNKNOWN FUNCTION

DOS 5.0

The purpose of this function has not been determined.

#### Call With

AX AD10h  
*additional arguments if any unknown*

#### Returns

CF clear if success  
 CF set on error

**Note** This function is a NOP if the keyboard package is not loaded. INT AD02h returns BX-FFFFh if its purpose otherwise is not known.

### INT 2Fh Function AD40h UNKNOWN FUNCTION

DOS 4+

The purpose of this function has not been determined.

#### Call With

AX AD40h  
 DX *unknown*  
*additional arguments if any unknown*

#### Returns

*unknown*

**Note** C:\DOS\3.31\DOS40.ZIP\EXECM\EXE\DOS3\_2\PRINTM

### INT 2Fh Function AD80h KEYB.COM INSTALLATION CHECK

DOS 3.3+

Determine whether KEYB.COM has been loaded.

#### Call With

AX AD80h

#### Returns

AI, HI if installed  
 BX version number (major in BH, minor in BL)  
 ENDI pointer to internal data (see below)

**Note** MS-DOS 3.30 [PC DOS 3.31] - [MS-DOS 5.00 & 6.00] use a value of 1

This function is also used to determine if the keyboard package is installed in unformatted disk.

#### Format of KEYB internal data.

Offset	Size	Description
00h	DWORD	original INT 09h
04h	DWORD	original INT 21h
08h	6 BYTES	<i>unknown</i>
0Eh	WORD	" "
10h	BYTE	" "
11h	BYTE	" "
12h	4 BYTES	<i>unknown</i>
16h	2 BYTES	country ID letters
18h	WORD	current code page

#### —DOS 3.3—

44: 4081 *private internal data, not documented*  
 4C: 408D *private internal data, not documented*

2 BYTES	unknown
WORD	pointer to key translation data
WORD	pointer to last 0x0000 code page table list (see below)

**DOS 4.01+**

2 BYTES	
WORD	
WORD	
2 BYTES	
WORD	pointer to key translation data
WORD	pointer to last item in code page table list (see below)
2 BYTES	unknown

**Format of code page table list entries**

Offset	Size	Description
00h	WORD	pointer to next item (FFFFh if last)
	WORD	code page
	BYTES	unknown

**Format of translation data**

Offset	Size	Description
00h	WORD	size of data in bytes (including this word)
	BYTES	unknown

**INT 2Fh Function AD81h****DOS 3.3+****KEYB COM SET KEYBOARD CODE PAGE**

Select a new code page list to use by the keyboard.

**Call With**

AX	AD81h
BX	code page (see 2F 0000h)

**Returns**

CF	set on error
AX	0000h (code page not changed)
CF	clear if OK

**Note** This function is not available in DOS 3.0, 3.1, 3.2, 3.3, 3.4, 3.5, 3.6, 3.7, 3.8, 3.9, 4.0, 4.1, 4.2, 4.3, 4.4, 4.5, 4.6, 4.7, 4.8, 4.9, 5.0, 5.1, 5.2, 5.3, 5.4, 5.5, 5.6, 5.7, 5.8, 5.9, 6.0, 6.1, 6.2, 6.3, 6.4, 6.5, 6.6, 6.7, 6.8, 6.9, 7.0, 7.1, 7.2, 7.3, 7.4, 7.5, 7.6, 7.7, 7.8, 7.9, 8.0, 8.1, 8.2, 8.3, 8.4, 8.5, 8.6, 8.7, 8.8, 8.9, 9.0, 9.1, 9.2, 9.3, 9.4, 9.5, 9.6, 9.7, 9.8, 9.9, 10.0, 10.1, 10.2, 10.3, 10.4, 10.5, 10.6, 10.7, 10.8, 10.9, 11.0, 11.1, 11.2, 11.3, 11.4, 11.5, 11.6, 11.7, 11.8, 11.9, 12.0, 12.1, 12.2, 12.3, 12.4, 12.5, 12.6, 12.7, 12.8, 12.9, 13.0, 13.1, 13.2, 13.3, 13.4, 13.5, 13.6, 13.7, 13.8, 13.9, 14.0, 14.1, 14.2, 14.3, 14.4, 14.5, 14.6, 14.7, 14.8, 14.9, 15.0, 15.1, 15.2, 15.3, 15.4, 15.5, 15.6, 15.7, 15.8, 15.9, 16.0, 16.1, 16.2, 16.3, 16.4, 16.5, 16.6, 16.7, 16.8, 16.9, 17.0, 17.1, 17.2, 17.3, 17.4, 17.5, 17.6, 17.7, 17.8, 17.9, 18.0, 18.1, 18.2, 18.3, 18.4, 18.5, 18.6, 18.7, 18.8, 18.9, 19.0, 19.1, 19.2, 19.3, 19.4, 19.5, 19.6, 19.7, 19.8, 19.9, 20.0, 20.1, 20.2, 20.3, 20.4, 20.5, 20.6, 20.7, 20.8, 20.9, 21.0, 21.1, 21.2, 21.3, 21.4, 21.5, 21.6, 21.7, 21.8, 21.9, 22.0, 22.1, 22.2, 22.3, 22.4, 22.5, 22.6, 22.7, 22.8, 22.9, 23.0, 23.1, 23.2, 23.3, 23.4, 23.5, 23.6, 23.7, 23.8, 23.9, 24.0, 24.1, 24.2, 24.3, 24.4, 24.5, 24.6, 24.7, 24.8, 24.9, 25.0, 25.1, 25.2, 25.3, 25.4, 25.5, 25.6, 25.7, 25.8, 25.9, 26.0, 26.1, 26.2, 26.3, 26.4, 26.5, 26.6, 26.7, 26.8, 26.9, 27.0, 27.1, 27.2, 27.3, 27.4, 27.5, 27.6, 27.7, 27.8, 27.9, 28.0, 28.1, 28.2, 28.3, 28.4, 28.5, 28.6, 28.7, 28.8, 28.9, 29.0, 29.1, 29.2, 29.3, 29.4, 29.5, 29.6, 29.7, 29.8, 29.9, 30.0, 30.1, 30.2, 30.3, 30.4, 30.5, 30.6, 30.7, 30.8, 30.9, 31.0, 31.1, 31.2, 31.3, 31.4, 31.5, 31.6, 31.7, 31.8, 31.9, 32.0, 32.1, 32.2, 32.3, 32.4, 32.5, 32.6, 32.7, 32.8, 32.9, 33.0, 33.1, 33.2, 33.3, 33.4, 33.5, 33.6, 33.7, 33.8, 33.9, 34.0, 34.1, 34.2, 34.3, 34.4, 34.5, 34.6, 34.7, 34.8, 34.9, 35.0, 35.1, 35.2, 35.3, 35.4, 35.5, 35.6, 35.7, 35.8, 35.9, 36.0, 36.1, 36.2, 36.3, 36.4, 36.5, 36.6, 36.7, 36.8, 36.9, 37.0, 37.1, 37.2, 37.3, 37.4, 37.5, 37.6, 37.7, 37.8, 37.9, 38.0, 38.1, 38.2, 38.3, 38.4, 38.5, 38.6, 38.7, 38.8, 38.9, 39.0, 39.1, 39.2, 39.3, 39.4, 39.5, 39.6, 39.7, 39.8, 39.9, 40.0, 40.1, 40.2, 40.3, 40.4, 40.5, 40.6, 40.7, 40.8, 40.9, 41.0, 41.1, 41.2, 41.3, 41.4, 41.5, 41.6, 41.7, 41.8, 41.9, 42.0, 42.1, 42.2, 42.3, 42.4, 42.5, 42.6, 42.7, 42.8, 42.9, 43.0, 43.1, 43.2, 43.3, 43.4, 43.5, 43.6, 43.7, 43.8, 43.9, 44.0, 44.1, 44.2, 44.3, 44.4, 44.5, 44.6, 44.7, 44.8, 44.9, 45.0, 45.1, 45.2, 45.3, 45.4, 45.5, 45.6, 45.7, 45.8, 45.9, 46.0, 46.1, 46.2, 46.3, 46.4, 46.5, 46.6, 46.7, 46.8, 46.9, 47.0, 47.1, 47.2, 47.3, 47.4, 47.5, 47.6, 47.7, 47.8, 47.9, 48.0, 48.1, 48.2, 48.3, 48.4, 48.5, 48.6, 48.7, 48.8, 48.9, 49.0, 49.1, 49.2, 49.3, 49.4, 49.5, 49.6, 49.7, 49.8, 49.9, 50.0, 50.1, 50.2, 50.3, 50.4, 50.5, 50.6, 50.7, 50.8, 50.9, 51.0, 51.1, 51.2, 51.3, 51.4, 51.5, 51.6, 51.7, 51.8, 51.9, 52.0, 52.1, 52.2, 52.3, 52.4, 52.5, 52.6, 52.7, 52.8, 52.9, 53.0, 53.1, 53.2, 53.3, 53.4, 53.5, 53.6, 53.7, 53.8, 53.9, 54.0, 54.1, 54.2, 54.3, 54.4, 54.5, 54.6, 54.7, 54.8, 54.9, 55.0, 55.1, 55.2, 55.3, 55.4, 55.5, 55.6, 55.7, 55.8, 55.9, 56.0, 56.1, 56.2, 56.3, 56.4, 56.5, 56.6, 56.7, 56.8, 56.9, 57.0, 57.1, 57.2, 57.3, 57.4, 57.5, 57.6, 57.7, 57.8, 57.9, 58.0, 58.1, 58.2, 58.3, 58.4, 58.5, 58.6, 58.7, 58.8, 58.9, 59.0, 59.1, 59.2, 59.3, 59.4, 59.5, 59.6, 59.7, 59.8, 59.9, 60.0, 60.1, 60.2, 60.3, 60.4, 60.5, 60.6, 60.7, 60.8, 60.9, 61.0, 61.1, 61.2, 61.3, 61.4, 61.5, 61.6, 61.7, 61.8, 61.9, 62.0, 62.1, 62.2, 62.3, 62.4, 62.5, 62.6, 62.7, 62.8, 62.9, 63.0, 63.1, 63.2, 63.3, 63.4, 63.5, 63.6, 63.7, 63.8, 63.9, 64.0, 64.1, 64.2, 64.3, 64.4, 64.5, 64.6, 64.7, 64.8, 64.9, 65.0, 65.1, 65.2, 65.3, 65.4, 65.5, 65.6, 65.7, 65.8, 65.9, 66.0, 66.1, 66.2, 66.3, 66.4, 66.5, 66.6, 66.7, 66.8, 66.9, 67.0, 67.1, 67.2, 67.3, 67.4, 67.5, 67.6, 67.7, 67.8, 67.9, 68.0, 68.1, 68.2, 68.3, 68.4, 68.5, 68.6, 68.7, 68.8, 68.9, 69.0, 69.1, 69.2, 69.3, 69.4, 69.5, 69.6, 69.7, 69.8, 69.9, 70.0, 70.1, 70.2, 70.3, 70.4, 70.5, 70.6, 70.7, 70.8, 70.9, 71.0, 71.1, 71.2, 71.3, 71.4, 71.5, 71.6, 71.7, 71.8, 71.9, 72.0, 72.1, 72.2, 72.3, 72.4, 72.5, 72.6, 72.7, 72.8, 72.9, 73.0, 73.1, 73.2, 73.3, 73.4, 73.5, 73.6, 73.7, 73.8, 73.9, 74.0, 74.1, 74.2, 74.3, 74.4, 74.5, 74.6, 74.7, 74.8, 74.9, 75.0, 75.1, 75.2, 75.3, 75.4, 75.5, 75.6, 75.7, 75.8, 75.9, 76.0, 76.1, 76.2, 76.3, 76.4, 76.5, 76.6, 76.7, 76.8, 76.9, 77.0, 77.1, 77.2, 77.3, 77.4, 77.5, 77.6, 77.7, 77.8, 77.9, 78.0, 78.1, 78.2, 78.3, 78.4, 78.5, 78.6, 78.7, 78.8, 78.9, 79.0, 79.1, 79.2, 79.3, 79.4, 79.5, 79.6, 79.7, 79.8, 79.9, 80.0, 80.1, 80.2, 80.3, 80.4, 80.5, 80.6, 80.7, 80.8, 80.9, 81.0, 81.1, 81.2, 81.3, 81.4, 81.5, 81.6, 81.7, 81.8, 81.9, 82.0, 82.1, 82.2, 82.3, 82.4, 82.5, 82.6, 82.7, 82.8, 82.9, 83.0, 83.1, 83.2, 83.3, 83.4, 83.5, 83.6, 83.7, 83.8, 83.9, 84.0, 84.1, 84.2, 84.3, 84.4, 84.5, 84.6, 84.7, 84.8, 84.9, 85.0, 85.1, 85.2, 85.3, 85.4, 85.5, 85.6, 85.7, 85.8, 85.9, 86.0, 86.1, 86.2, 86.3, 86.4, 86.5, 86.6, 86.7, 86.8, 86.9, 87.0, 87.1, 87.2, 87.3, 87.4, 87.5, 87.6, 87.7, 87.8, 87.9, 88.0, 88.1, 88.2, 88.3, 88.4, 88.5, 88.6, 88.7, 88.8, 88.9, 89.0, 89.1, 89.2, 89.3, 89.4, 89.5, 89.6, 89.7, 89.8, 89.9, 90.0, 90.1, 90.2, 90.3, 90.4, 90.5, 90.6, 90.7, 90.8, 90.9, 91.0, 91.1, 91.2, 91.3, 91.4, 91.5, 91.6, 91.7, 91.8, 91.9, 92.0, 92.1, 92.2, 92.3, 92.4, 92.5, 92.6, 92.7, 92.8, 92.9, 93.0, 93.1, 93.2, 93.3, 93.4, 93.5, 93.6, 93.7, 93.8, 93.9, 94.0, 94.1, 94.2, 94.3, 94.4, 94.5, 94.6, 94.7, 94.8, 94.9, 95.0, 95.1, 95.2, 95.3, 95.4, 95.5, 95.6, 95.7, 95.8, 95.9, 96.0, 96.1, 96.2, 96.3, 96.4, 96.5, 96.6, 96.7, 96.8, 96.9, 97.0, 97.1, 97.2, 97.3, 97.4, 97.5, 97.6, 97.7, 97.8, 97.9, 98.0, 98.1, 98.2, 98.3, 98.4, 98.5, 98.6, 98.7, 98.8, 98.9, 99.0, 99.1, 99.2, 99.3, 99.4, 99.5, 99.6, 99.7, 99.8, 99.9, 100.0, 100.1, 100.2, 100.3, 100.4, 100.5, 100.6, 100.7, 100.8, 100.9, 101.0, 101.1, 101.2, 101.3, 101.4, 101.5, 101.6, 101.7, 101.8, 101.9, 102.0, 102.1, 102.2, 102.3, 102.4, 102.5, 102.6, 102.7, 102.8, 102.9, 103.0, 103.1, 103.2, 103.3, 103.4, 103.5, 103.6, 103.7, 103.8, 103.9, 104.0, 104.1, 104.2, 104.3, 104.4, 104.5, 104.6, 104.7, 104.8, 104.9, 105.0, 105.1, 105.2, 105.3, 105.4, 105.5, 105.6, 105.7, 105.8, 105.9, 106.0, 106.1, 106.2, 106.3, 106.4, 106.5, 106.6, 106.7, 106.8, 106.9, 107.0, 107.1, 107.2, 107.3, 107.4, 107.5, 107.6, 107.7, 107.8, 107.9, 108.0, 108.1, 108.2, 108.3, 108.4, 108.5, 108.6, 108.7, 108.8, 108.9, 109.0, 109.1, 109.2, 109.3, 109.4, 109.5, 109.6, 109.7, 109.8, 109.9, 110.0, 110.1, 110.2, 110.3, 110.4, 110.5, 110.6, 110.7, 110.8, 110.9, 111.0, 111.1, 111.2, 111.3, 111.4, 111.5, 111.6, 111.7, 111.8, 111.9, 112.0, 112.1, 112.2, 112.3, 112.4, 112.5, 112.6, 112.7, 112.8, 112.9, 113.0, 113.1, 113.2, 113.3, 113.4, 113.5, 113.6, 113.7, 113.8, 113.9, 114.0, 114.1, 114.2, 114.3, 114.4, 114.5, 114.6, 114.7, 114.8, 114.9, 115.0, 115.1, 115.2, 115.3, 115.4, 115.5, 115.6, 115.7, 115.8, 115.9, 116.0, 116.1, 116.2, 116.3, 116.4, 116.5, 116.6, 116.7, 116.8, 116.9, 117.0, 117.1, 117.2, 117.3, 117.4, 117.5, 117.6, 117.7, 117.8, 117.9, 118.0, 118.1, 118.2, 118.3, 118.4, 118.5, 118.6, 118.7, 118.8, 118.9, 119.0, 119.1, 119.2, 119.3, 119.4, 119.5, 119.6, 119.7, 119.8, 119.9, 120.0, 120.1, 120.2, 120.3, 120.4, 120.5, 120.6, 120.7, 120.8, 120.9, 121.0, 121.1, 121.2, 121.3, 121.4, 121.5, 121.6, 121.7, 121.8, 121.9, 122.0, 122.1, 122.2, 122.3, 122.4, 122.5, 122.6, 122.7, 122.8, 122.9, 123.0, 123.1, 123.2, 123.3, 123.4, 123.5, 123.6, 123.7, 123.8, 123.9, 124.0, 124.1, 124.2, 124.3, 124.4, 124.5, 124.6, 124.7, 124.8, 124.9, 125.0, 125.1, 125.2, 125.3, 125.4, 125.5, 125.6, 125.7, 125.8, 125.9, 126.0, 126.1, 126.2, 126.3, 126.4, 126.5, 126.6, 126.7, 126.8, 126.9, 127.0, 127.1, 127.2, 127.3, 127.4, 127.5, 127.6, 127.7, 127.8, 127.9, 128.0, 128.1, 128.2, 128.3, 128.4, 128.5, 128.6, 128.7, 128.8, 128.9, 129.0, 129.1, 129.2, 129.3, 129.4, 129.5, 129.6, 129.7, 129.8, 129.9, 130.0, 130.1, 130.2, 130.3, 130.4, 130.5, 130.6, 130.7, 130.8, 130.9, 131.0, 131.1, 131.2, 131.3, 131.4, 131.5, 131.6, 131.7, 131.8, 131.9, 132.0, 132.1, 132.2, 132.3, 132.4, 132.5, 132.6, 132.7, 132.8, 132.9, 133.0, 133.1, 133.2, 133.3, 133.4, 133.5, 133.6, 133.7, 133.8, 133.9, 134.0, 134.1, 134.2, 134.3, 134.4, 134.5, 134.6, 134.7, 134.8, 134.9, 135.0, 135.1, 135.2, 135.3, 135.4, 135.5, 135.6, 135.7, 135.8, 135.9, 136.0, 136.1, 136.2, 136.3, 136.4, 136.5, 136.6, 136.7, 136.8, 136.9, 137.0, 137.1, 137.2, 137.3, 137.4, 137.5, 137.6, 137.7, 137.8, 137.9, 138.0, 138.1, 138.2, 138.3, 138.4, 138.5, 138.6, 138.7, 138.8, 138.9, 139.0, 139.1, 139.2, 139.3, 139.4, 139.5, 139.6, 139.7, 139.8, 139.9, 140.0, 140.1, 140.2, 140.3, 140.4, 140.5, 140.6, 140.7, 140.8, 140.9, 141.0, 141.1, 141.2, 141.3, 141.4, 141.5, 141.6, 141.7, 141.8, 141.9, 142.0, 142.1, 142.2, 142.3, 142.4, 142.5, 142.6, 142.7, 142.8, 142.9, 143.0, 143.1, 143.2, 143.3, 143.4, 143.5, 143.6, 143.7, 143.8, 143.9, 144.0, 144.1, 144.2, 144.3, 144.4, 144.5, 144.6, 144.7, 144.8, 144.9, 145.0, 145.1, 145.2, 145.3, 145.4, 145.5, 145.6, 145.7, 145.8, 145.9, 146.0, 146.1, 146.2, 146.3, 146.4, 146.5, 146.6, 146.7, 146.8, 146.9, 147.0, 147.1, 147.2, 147.3, 147.4, 147.5, 147.6, 147.7, 147.8, 147.9, 148.0, 148.1, 148.2, 148.3, 148.4, 148.5, 148.6, 148.7, 148.8, 148.9, 149.0, 149.1, 149.2, 149.3, 149.4, 149.5, 149.6, 149.7, 149.8, 149.9, 150.0, 150.1, 150.2, 150.3, 150.4, 150.5, 150.6, 150.7, 150.8, 150.9, 151.0, 151.1, 151.2, 151.3, 151.4, 151.5, 151.6, 151.7, 151.8, 151.9, 152.0, 152.1, 152.2, 152.3, 152.4, 152.5, 152.6, 152.7, 152.8, 152.9, 153.0, 153.1, 153.2, 153.3, 153.4, 153.5, 153.6, 153.7, 153.8, 153.9, 154.0, 154.1, 154.2, 154.3, 154.4, 154.5, 154.6, 154.7, 154.8, 154.9, 155.0, 155.1, 155.2, 155.3, 155.4, 155.5, 155.6, 155.7, 155.8, 155.9, 156.0, 156.1, 156.2, 156.3, 156.4, 156.5, 156.6, 156.7, 156.8, 156.9, 157.0, 157.1, 157.2, 157.3, 157.4, 157.5, 157.6, 157.7, 157.8, 157.9, 158.0, 158.1, 158.2, 158.3, 158.4, 158.5, 158.6, 158.7, 158.8, 158.9, 159.0, 159.1, 159.2, 159.3, 159.4, 159.5, 159.6, 159.7, 159.8, 159.9, 160.0, 160.1, 160.2, 160.3, 160.4, 160.5, 160.6, 160.7, 160.8, 160.9, 161.0, 161.1, 161.2, 161.3, 161.4, 161.5, 161.6, 161.7, 161.8, 161.9, 162.0, 162.1, 162.2, 162.3, 162.4, 162.5, 162.6, 162.7, 162.8, 162.9, 163.0, 163.1, 163.2, 163.3, 163.4, 163.5, 163.6, 163.7, 163.8, 163.9, 164.0, 164.1, 164.2, 164.3, 164.4, 164.5, 164.6, 164.7, 164.8, 164.9, 165.0, 165.1, 165.2, 165.3, 165.4, 165.5, 165.6, 165.7, 165.8, 165.9, 166.0, 166.1, 166.2, 166.3, 166.4, 166.5, 166.6, 166.7, 166.8, 166.9, 167.0, 167.1, 167.2, 167.3, 167.4, 167.5, 167.6, 167.7, 167.8, 167.9, 168.0, 168.1, 168.2, 168.3, 168.4, 168.5, 168.6, 168.7, 168.8, 168.9, 169.0, 169.1, 169.2, 169.3, 169.4, 169.5, 169.6, 169.7, 169.8, 169.9, 170.0, 170.1, 170.2, 170.3, 170.4, 170.5, 170.6, 170.7, 170.8, 170.9, 171.0, 171.1, 171.2, 171.3, 171.4, 171.5, 171.6, 171.7, 171.8, 171.9, 172.0, 172.1, 172.2, 172.3, 172.4, 172.5, 172.6, 172.7, 172.8, 172.9, 173.0, 173.1, 173.2, 173.3, 173.4, 173.5, 173.6, 173.7, 173.8, 173.9, 174.0, 174.1, 174.2, 174.3, 174.4, 174.5, 174.6, 174.7, 174.8, 174.9, 175.0, 175.1, 175.2, 175.3, 175.4, 175.5, 175.6, 175.7, 175.8, 175.9, 1



**INT 2Fh Function AD83h****DOS 5+****KEYB COM - GET KEYBOARD MAPPING**

Executes only if the keyboard driver is using the standard keyboard type (code 0001).

Called With

AX A1083h

Returns

BI current status  
 00h US keyboard  
 FFh foreign keyboard

See Also: 2F A1082h

**INT 2Fh Function AE00h****DOS 3.3+****INSTALLABLE COMMAND INSTALLATION CHECK**

Executes only if the command driver is INR extension to COMMAND.COM's command set

Called With

AX A100Fh  
 DX FFFFh  
 CX FFh first call 0-second call  
 C[ length of command line call 4100F40  
 DS:BX pointer to command line buffer see below  
 DS:SI pointer to command name buffer see below  
 DI (0000) 4100F40

Returns

SI pointer to COMMAND.COM  
 AI 00h if the command should be executed as usual

**Notes:** This call provides a means for installing new commands into the command set of a command driver. The command driver must be INR extension to COMMAND.COM. The command driver must be loaded before the command set is initialized. The command driver must be loaded before the command set is initialized.

See APPENDIX

See Also: INT 2Fh

**Format of command line buffer**

Offset	Size	Description
00h	BYTE	max length of command line as in 21 0Ah
01h	BYTE	max length of command line as in 21 0Ah
02h	N BYTES	command line text terminated by 00h

**Format of command name buffer**

Offset	Size	Description
00h	BYTE	length of command line
01h	N BYTES	command line text terminated by 00h

**INT 2Fh Function AE01h****DOS 3.3+****INSTALLABLE COMMAND EXECUTE**

Execute a INR extension to COMMAND.COM's set of internal commands. The extension may require the use of the command driver.

Called With

AX AE01h  
 DX FFFFh



**INT 2Fh Function B701h****APPEND****GET APPEND PATH**

Alternative to function B704h to retrieve the current active directory path. Each APPEND will search for a file.

Call With

AX B701h

Returns

ES:DI pointer to active APPEND path

**Note:** Not all versions of APPEND support this call. See AX-B704h for and only call this function if that one is not supported.

MIN DOS 3.30 APPEND; incorrect APPEND version and obsolete call.

See Also: AX-B704h

**INT 2Fh Function B702h****DOS 3.3+****APPEND VERSION CHECK**

Determine which version of APPEND has been loaded.

Call With

AX B702h

Returns

EH:ESI not DOS 4.0 APPEND; also DOS 8.0 APPEND

AX major version number

AH minor version number, otherwise

**Note:** This function is undocumented for DOS 5.0.

See Also: 2-F710h

**INT 2Fh Function B703h****DOS 3.3, 5.0****APPEND HOOK INT 21h**

Specifies the handler APPEND should use when the file has finished processing. (min DOS 3.31)

Call With

AX B703h

ES:DI pointer to INT 21h under APPEND should call to

Returns

ESI pointer to APPEND's INT 21h handler

**Note:** Each invocation of this function implies a flag which APPEND uses to determine whether to chain to the user handler or the original INT 21h.

**INT 2Fh Function B704h****DOS 3.3+****APPEND GET APPEND PATH**

Return the current APPEND path.

Call With

AX B704h

Returns

ES:DI pointer to active APPEND path (28 bytes max)

**Note:** This function is documented for DOS 5.0. Some versions of append do not support this call, and return. See also call. In this case, you should call AX-B701h to get the APPEND path.

See Also: AX-B701h

**INT 2Fh Function B706h****DOS 4+****APPEND GET APPEND FUNCTION STATE**

Determines which actions APPEND is performing.

Call With

AX B706h

Returns

BX APPEND state bits, see Table B706h  
 bit 0 set: CALLPND set  
 bit 1 set: reserved  
 bit 2: DOS50: set if APPEND applies directory search even if a  
     time has been specified  
 bit 3 set: PATH flag active  
 bit 4 set: E flag active, environment var APPEND exists  
 bit 5 set: A flag

**INT 2Fh Function B707h****DOS 4+****APPEND SET APPEND FUNCTION STATE**

Specifies which actions APPEND supports.

Call With

AX B706h  
 BX APPEND state bits, see Table B706h

**INT 2Fh Function B710h****DOS 3.3+****APPEND GET VERSION INFO**

Determines a field version of APPEND (see Appendix B).

Call With

AX

Returns

AX version of MESS (state, see Table B706h)  
 BX major version (0000-03) (00 = standard version)  
 CX minor version (0000-03)  
 DI 0000  
 DI1 0

See Also: Table B706h

**INT 2Fh Function B711h****DOS 4+****APPEND SET RETURN FOUND NAME STATE**

Specifies that the file specified by FILENAME be written over the filename passed to the next INT 2Fh call.

Call With

AX B711h  
 BX FILENAME  
 CX FILENAME  
 DI FILENAME  
 DI1 FILENAME

2Fh call processed by APPEND

BUG: DOS 4.01: END reverts to overwrites DS:DI instead of DS:SI for 2Fh 6Ch

See Also: Table B711h

**INT 2Fh Function 8800h****Network****INSTALLATION CHECK**

Determines whether a network server

## Call With

AX B800h

## Returns

AL status  
 00h not installed  
 non-zero installed

Bit 0 installed/computer flags (set in this order)

bit 0 network

bit 2 message

bit 7 receive

bit 3 redirector

**INT 2Fh Function B803h****Network****GET NETWORK EVENT POST HANDLER**

This function is used in conjunction with 2F-B804h to hook into the network event post routine.

## Call With

AX B803h

## Returns

ES:BX pointer to current post handler (see 2F-B804h)

See Also: 2F-B804h, 2F-B803h

**INT 2Fh Function B804h****Network****SET NETWORK EVENT POST HANDLER**

This function is used in conjunction with 2F-B803h to hook into the network event post routine.

## Call With

AX B804h

ES:BX pointer to new event post handler

**Note:** DS:SI, DS:DI, DS:SI are called for in network event post routine to detect message received and critical network error.

See Also: 2F-B803h, 2F-B804h

**Values post routine is called with:**

AX 0000h end of block routine

DS:SI pointer to ASCII originator name

DS:DI pointer to ASCII destination name

ES:BX pointer to text header (see below)

AX 0001h start multiple block message

CX block group ID

DS:SI pointer to ASCII originator name

DS:DI pointer to ASCII destination name

AX 0002h multiple block text

CX block group ID

ES:BX pointer to text header (see below)

AX 0003h end multiple block message

CX block group ID

AX 0004h message aborted due to error

CX block group ID

AX 0005h server received badly formatted network request

Return:

AX FFFh:PC LAN will process error

AX 0102h unexpected network error



---

**INT 2Fh Function B809h** **Network PC**  
**LAN Program, LAN Manager, DOS LAN Requester Version Check**

Determine which version of the network software has been installed. Note that these networks return the version number in the way LANtastic and NetWare Lite return the version.

Call With

AX = 809h

Returns

AX = minor version (decimal)

AL = major version

See Also: AX-4F53h AX-B800h

---

**INT 2Fh Function B900h** **PC Network**  
**RECEIVER.COM INSTALLATION CHECK**

Determine if RECEIVER.COM is installed on the system.

Call With

AX = B900h

Returns

AL = 00h if not installed

FFh if installed

---

**INT 2Fh Function B901h** **PC Network**  
**RECEIVER.COM GET RECEIVER.COM INT 2Fh HANDLER ADDRESS**

Determine the entry point to a RECEIVER.COM INT 2Fh handler. See [RECEIVER.COM](#) for details on how to call this by bypassing the normal GET RECEIVER.COM INT 2Fh handler.

Call With

AX = B901h

Returns

AX = unknown

ES:BX = pointer to RECEIVER.COM INT 2Fh handler

---

**INT 2Fh Function B903h** **PC Network**  
**RECEIVER.COM GET RECEIVER.COM POST ADDRESS**

Find the entry point to a RECEIVER.COM POST handler.

Call With

AX = B903h

Returns

ES:AX = pointer to POST handler

See Also: 2F-B803h 2F-B904h

---

**INT 2Fh Function B904h** **PC Network**  
**RECEIVER.COM SET RECEIVER.COM POST ADDRESS**

Find the entry point to a RECEIVER.COM POST handler. See [RECEIVER.COM](#) for details.

Call With

AX = B904h

ES:BX = pointer to new POST handler

See Also: 2F-B804h 2F-B905h

---

**INT 2Fh Function B905h****PC Network****RECEIVER COM GET FILENAME**

Determines the filename used internally by RECEIVER.COM

Call With

- AX B905h
- DS:BX pointer to 128 byte buffer for filename 1
- ES:DX pointer to 128 byte buffer for filename 2

Returns

00h if 1 from RECEIVER.COM internal; 01h if 2

See Also

B.34265

See Also: 2F.1506h

**INT 2Fh Function B906h****PC Network****RECEIVER COM SET FILENAME**

Specifies the filename which RECEIVER.COM uses internally

Call With

- AX B906h
- DS:BX pointer to 128 byte buffer for filename 1
- ES:DX pointer to 128 byte buffer for filename 2

Returns

00h if RECEIVER.COM accepted; 01h if failed from user buffer

See Also

B.34265, B.34266

See Also: 2F.1506h

**INT 2Fh Function B908h****PC Network****RECEIVER COM UNLINK KEYBOARD HANDLER**

CALLS THE INT 10h BIOS

Call With

- AX B908h

See Also

B.34199

B.34200

B.34201

B.34202

B.34203

B.34204

B.34205

B.34206

B.34207

B.34208

B.34209

B.34210

B.34211

B.34212

B.34213

B.34214

B.34215

B.34216

B.34217

B.34218

B.34219

B.34220

B.34221

B.34222

B.34223

B.34224

B.34225

B.34226

B.34227

B.34228

B.34229

B.34230

B.34231

B.34232

B.34233

B.34234

B.34235

B.34236

B.34237

B.34238

B.34239

B.34240

B.34241

B.34242

B.34243

B.34244

B.34245

B.34246

B.34247

B.34248

B.34249

B.34250

B.34251

B.34252

B.34253

B.34254

B.34255

B.34256

B.34257

B.34258

B.34259

B.34260

B.34261

B.34262

B.34263

B.34264

B.34265

B.34266

B.34267

B.34268

B.34269

B.34270

B.34271

B.34272

B.34273

B.34274

B.34275

B.34276

B.34277

B.34278

B.34279

B.34280

B.34281

B.34282

B.34283

B.34284

B.34285

B.34286

B.34287

B.34288

B.34289

B.34290

B.34291

B.34292

B.34293

B.34294

B.34295

B.34296

B.34297

B.34298

B.34299

B.34300

B.34301

B.34302

B.34303

B.34304

B.34305

B.34306

B.34307

B.34308

B.34309

B.34310

B.34311

B.34312

B.34313

B.34314

B.34315

B.34316

B.34317

B.34318

B.34319

B.34320

B.34321

B.34322

B.34323

B.34324

B.34325

B.34326

B.34327

B.34328

B.34329

B.34330

B.34331

B.34332

B.34333

B.34334

B.34335

B.34336

B.34337

B.34338

B.34339

B.34340

B.34341

B.34342

B.34343

B.34344

B.34345

B.34346

B.34347

B.34348

B.34349

B.34350

B.34351

B.34352

B.34353

B.34354

B.34355

B.34356

B.34357

B.34358

B.34359

B.34360

B.34361

B.34362

B.34363

B.34364

B.34365

B.34366

B.34367

B.34368

B.34369

B.34370

B.34371

B.34372

B.34373

B.34374

B.34375

B.34376

B.34377

B.34378

B.34379

B.34380

B.34381

B.34382

B.34383

B.34384

B.34385

B.34386

B.34387

B.34388

B.34389

B.34390

B.34391

B.34392

B.34393

B.34394

B.34395

B.34396

B.34397

B.34398

B.34399

B.34400

B.34401

B.34402

B.34403

B.34404

B.34405

B.34406

B.34407

B.34408

B.34409

B.34410

B.34411

B.34412

B.34413

B.34414

B.34415

B.34416

B.34417

B.34418

B.34419

B.34420

B.34421

B.34422

B.34423

B.34424

B.34425

B.34426

B.34427

B.34428

B.34429

B.34430

B.34431

B.34432

B.34433

B.34434

B.34435

B.34436

B.34437

B.34438

B.34439

B.34440

B.34441

B.34442

B.34443

B.34444

B.34445

B.34446

B.34447

B.34448

B.34449

B.34450

B.34451

B.34452

B.34453

B.34454

B.34455

B.34456

B.34457

B.34458

B.34459

B.34460

B.34461

B.34462

B.34463

B.34464

B.34465

B.34466

B.34467

B.34468

B.34469

B.34470

B.34471

B.34472

B.34473

B.34474

B.34475

B.34476

B.34477

B.34478

B.34479

B.34480

B.34481

B.34482

B.



**INT 2Fh Function BC06h****Windows 3.0, DOS 5+****EGA.SYS GET VERSION INFO**

Determine which version of EGA.SYS has been loaded.

Call With

AX = BC 06h

Returns

BX = 5456h "13"

CH = major version

CL = minor version

DI = revision

See Also: *INT 10 AH-E, 16, 2E-BC 00h***INT 2Fh Function BF00h****PC LAN****REDIRFS.EXE INSTALLATION CHECK**

Determine whether the PC LAN Program has been installed successfully.

Call With

AX = BF 00h

Returns

AL = 1Fh if installed

**INT 2Fh Function BF01h****PC LAN****REDIRFS.EXE UNKNOWN FUNCTION**

The purpose of this function has not been determined.

Call With

AX = BF 01h

*additional arguments if any unknown*

Returns

*unknown***INT 2Fh Function BF80h****PC LAN****REDIR.SYS SET REDIRFS ENTRY POINT**

Specify the address of the RedirFS workspace for the PC LAN program.

Call With

AX = BF 80h

ESI = DS:EDI = ES:EDI = FS:EDI = IP:IP2 = IP:IP3 = IP:IP4 = IP:IP5

Returns

AX = 0000h

FS:EDI pointer to internal workspace

Note: All other registers must be zero. Note: REDIRFS.SYS is a DOS 5.03-specific extension.

**INT 30h****DOS 1+****FAR JMP instruction for CP/M-style calls**

This instruction is used to call DOS 1.0 through 2.22, CP/M 1.0 through 2.22, and CP/M 2.0 through 2.22. The instruction is not used to call DOS 3.0 through 3.21, CP/M 2.23 through 2.24, or DOS 4.0 through 4.08. The instruction is not used to call DOS 5.0 through 5.03, CP/M 3.0 through 3.03, or DOS 6.0 through 6.22.

CALL instruction: The instruction is used to call DOS 1.0 through 2.22, CP/M 1.0 through 2.22, and CP/M 2.0 through 2.22. The instruction is not used to call DOS 3.0 through 3.21, CP/M 2.23 through 2.24, or DOS 4.0 through 4.08. The instruction is not used to call DOS 5.0 through 5.03, CP/M 3.0 through 3.03, or DOS 6.0 through 6.22.

JMP instruction: The instruction is used to call DOS 3.0 through 3.21, CP/M 2.23 through 2.24, or DOS 4.0 through 4.08. The instruction is not used to call DOS 1.0 through 2.22, CP/M 1.0 through 2.22, and CP/M 2.0 through 2.22. The instruction is not used to call DOS 5.0 through 5.03, CP/M 3.0 through 3.03, or DOS 6.0 through 6.22.

was created on the FATC call 71 40h  
See Also 71 26h

**INT 31h****DOS 1+****OVERWRITTEN BY CP/M JUMP INSTRUCTION IN INT 30h**

Call With AX = 0000h  
Returns AX = 0000h

detected in the list on the accompanying disk

Notes: All functions which are undocumented  
All functions which are undocumented

0051 Control selector  
MS-DOS program candidate  
downward pages

For more information see the interrupt list in *dosk* and *Endwin* & *Windows*, chapter 4

**INT 67h Function FF5Ah  
EMM386 INSTALLATION CHECK****Microsoft EMM386 EXE v4.20+**

Call With

AX = 125Ah

Returns

AX = 845Ah if loaded

Notes: AX pointer to API entry point

is the address of the EMM386 installation routine  
is the address of the EMM386 installation routine  
is the address of the EMM386 installation routine

**Call API entry point with:**

API = 0000h

Returns

API status

00h not active (00h)

01h "Active" mode

API = 0001h: memory management state

00h (00h 01h 01h) = 07h (01h 10h)

API = 0002h: Waitck coprocessor support

API = 0003h

00h: Waitck support state

Returns

API status

00h: Waitck coprocessor is present

01h: Waitck support is enabled

02h: turn on Waitck support

03h: turn off Waitck support

**— v4.20 4.41 only —**

API = 0004h: 00h: Windows support

API = 0005h: 00h: 01h

API = 0006h: 00h: 01h

API = 0007h: 00h: 01h

See 71 09h

MMIO on the DOS prompt

## Glossary

The following list is a partial and somewhat random list of some terms and abbreviations used throughout this book. *Accented terms* means to “see also” the italicized term. XX/YY and XX/YYY means to see the entry for INT XXh, INT YYh or XX/YYYh in the appendix on Undocumented DOS Functions and Data Structures.

**A20** On 80286 and higher processors, the A20 address line controls access to the first megabyte of extended memory (address 1–20 1000000). For compatibility with “address wraparound” on the older 8088 processor, 286+ processors leave A20 disabled. However, to access the *HMA* needed for DOS *HIGH*, A20 must be enabled. See chapter 6.

**ASCIIZ** A zero-terminated ASCII string, such as A B C D00.

**BPB** The BIOS Parameter Block stores the low-level layout of a drive. See 21-53 and chapter 8.

**CDS** The Current Directory Structure for a drive stores the current directory, type, and other information about a logical drive. See 21-52 and chapter 8.

**Chicago** Microsoft's forthcoming desktop operating system, incorporating DOS 7 and Windows 4, includes *protected mode*, *multitasking*, and other features. See 21-4302, 21/71, 21/72, and chapter 8.

**Conventional memory** Memory below one megabyte (1000000) immediately accessible by a *real mode* program. Contrast *extended memory*.

**DOS extender** Software that provides INT 21h and other DOS services in protected mode, creating the illusion that MS-DOS is a protected mode operating system. Windows contains a DOS extender. See chapter 3.

**DOSMGU** The component of Windows Enhanced mode responsible for interfacing with MS-DOS. DOSMGU also provides the Enhanced mode *DOS extender*. See 21-1667, chapter 4, and chapter 3.

**DPB** The DOS Device Parameter Block stores the description of the device layout for a logical drive, as well as some housekeeping information. See 21-11-21-32 and chapter 8.

**DblSpace** *DblSpace* provides on-the-fly disk compression in MS-DOS 6.0 and higher. See 21/4A11 and chapter 8.

**DPL** The DOS Parameter List is used to pass arguments to SHARE and network functions. See 21/5100.

**DPMI** DOS programs can use the DOS Protected Mode Interface to switch themselves into protected mode. Once in protected mode, they can use DPMI (INT 31h) services to communicate back to real mode. In Windows Enhanced mode, DPMI services are provided by the *EMM*. See chapter 3.

**DR DOS** DR DOS 5 and 6 were DOS workalikes produced by Digital Research, later bought by Novell. DR DOS has been renamed Novell DOS 7. See 21-4452 and chapter 4.

**DIA** The Disk Transfer Address indicates where functions which do not take an explicit data address will read or store data. Although the name implies that only disk accesses use this address, other functions use it as well. See 21-41 for the `FindNext` DIA format.

**Extended memory** Memory above one megabyte (100000h) usually accessible only from a *protected mode* program. XMS provides services that allow a real mode program to access extended memory.

**FAT** The File Allocation Table of a disk, which records the clusters that are in use. See chapter 8.

**FCB** A File Control Block, which is used by DOS 1.x functions to record the state of an open file. See 21-13. Since an FCB resides in an application's own address space, DOS maintains internal System FCBs (so-called SFCBs). See 21-52 and chapter 8.

**File handle** An index into a *JFF*.

**HMA** The High Memory Area is a slice of extended memory, just under 64k bytes, between addresses 00000 and 11110h + 1111h = 10011h, that is accessible from real mode on 80286 and higher processors. Real mode programs access the HMA with an address segment of 1111h. The HMA is not accessible if the 420 address line is enabled. See 21-4401 and 21-4402.

**IFS** An Interfile File System, which allows non-DOS format media to be used by DOS. In most ways, IFS is very similar to a networked drive, although an IFS would typically be local rather than remote. The original IFS interface was for DOS 4.x only. DOS 3, DOS 5, and DOS 6 use 21-11 for the network redirector interface. *Chirrup* will have a documented IFSMGR interface.

**JFF** The File File Table, also called Operation Table, usually stored in a program's PSP, which translates file handles to SFCB indices. See 21-26 and chapter 8.

**LoL** List of Lists—see *SeeVars*.

**MCB** A Memory Control Block (or ARMSA) contains the size and owner of a conventional memory allocation. See 21/52, and chapter 7.

**MZ** Headers of Mark Zdiskette (5AD3h) used as a signature for executable files. DOS also uses the letters M and Z as signatures for *MCB*.

**NCB** A Network Control Block used to pass requests to NETBIOS and receive status information from the NETBIOS handler. See NETSMB for the interrupt list on disk.

**NETX** Generic name for NetWare workstation shell (NE14EXE, NE14EXE, NE15EXE, etc.), which takes over the INT 21h interface. See chapter 4.

**Protected mode** Execution mode (real 32, 80286 and higher microprocessors), protected mode (real 32, 386, etc.) or virtual addressable access to memory above one megabyte. However, only 386+ microprocessors spend much of their time in real mode. See chapter 3.

**PSP** The Program Segment Prefix is a 256 byte data area prepended to a program when it is loaded. It contains the command line by which the program was invoked, a pointer to the *JFF* (and generally the `DI` field—the segment address of the program's environment), and a variety of housekeeping information for DOS. See also 21-26 and chapters 7 and 10.

**Real mode** The native mode of the Intel 8086 and 8086 microprocessors, limited to one megabyte of immediately addressable memory. 80286 and higher machines can emulate real mode, and in fact spend much of their time doing so. MS-DOS is a real mode operating system. See chapter 3.

**SDA** The DOS Swappable Data Area, containing many (though not all) of the variables used internally by DOS to record the state of a function call in progress. See 21-5D06, 21-5D0B, and chapters 8 and 9.

**SFT** A System File Table maintains the state of an open file for the DOS 2+ handle functions, just as an PCB maintains the state for DOS 1.x functions. See 21-5E2 and chapter 8. SFT file entries are also used for System PCBs.

**SysVars** Also known as the "List of Lists," List SysVars is a DOS internal data structure containing pointers to many other DOS internal structures, including the CDS, SFL, and so on. See 21-5F2.

**TSR** A program that calls INT 27h (Terminate and Stay Resident) or 21-3E (Keep Program). See chapter 9.

**UMB** An Upper Memory Block is an MB of that resides in a PC's "upper memory," the area between address A0000h and FFFFFh normally reserved for adapters. Expanded memory managers, including EMM386 in DOS 5 and higher, can allocate UMBS for use by DOS software. Unlike the HEMA, UMBS are always accessible in real mode.

**V86, VM, VMM** Virtual 8086 mode is provided by Intel 80386 and higher processors to emulate one or more 8088 like sessions. Each such session is called a Virtual Machine (VM). VMM can also include a protected mode component. 8086 memory managers such as EMM386, QEMM, and 386MAX, and multitaskers such as Windows Enhanced mode, run MS-DOS in V86 mode. All interrupts by a real mode software running in V86 mode are handled by a 62-bit protected mode program called a Virtual Machine Monitor (VMM). The VMM can emulate the interrupt function and can intercept the interrupt back to V86 mode. In Windows Enhanced mode, the VMM is contained inside WIN386.EXE, and works in conjunction with VxDs. For Chicago, the VMM is contained inside DOS386.EXE. See chapters 2, 3, and 4.

**VxD** Virtual Device Drivers are used in Windows Enhanced mode to emulate the behavior of and/or provide a multitasking interface to real mode software such as MS-DOS and the BIOS and devices such as the keyboard, display, disk, mouse, printer, and so on. See 21-1607, 21-16B4, and chapters 1 and 3. For OS 2.2.x and Windows NT, Virtual Device Drivers are called VDDs. See chapter 4.



## ANNOTATED BIBLIOGRAPHY

- Philip M. Adams and Clowis J. Londo: *Writing DOS Device Drivers in C*. Englewood Cliffs, NJ: Prentice Hall, 1990, 385 pp.
- Walter Adams and James Brock: *Antitrust Economics on Trial: A Dialogue on the New Laissez-Faire*. Project on NE. Princeton University Press, 1991, 152 pp. A brief, convincing polemic in the form of an *in vivo* against the Chicago School of Economics' defence of tying arrangements. A study of antipredatory practices discussed in chapter 1 of *Undocumented DOS*. Pages 30-37 in "the predation problem" are particularly useful.
- Alfred Aho, Brian Kernighan, and Peter Weinberger: *The AWK Programming Language*. Reading, MA: Addison-Wesley, 1988, 210 pp. AWK was an excellent language with which to build disassemblers and reverse engineering utilities.
- Prabhat K. Andhige: *UNIX System Utilities*. Englewood Cliffs, NJ: Prentice Hall, 1990, 274 pp. Every programmer interested in operating systems should know something about UNIX utilities. This is a good introduction, with very readable pseudocode. It is a good introduction when trying to tackle Bach's book (see below).
- Rick Aron: "DOS 6 Utilities vs. the Competition: What's the Price of a Free Lunch?" *PC Magazine*, September 14, 1993. On Microsoft's "bid to dominate the utilities world."
- Manuel J. Bach: *The Design of the UNIX Operating System*. Englewood Cliffs, NJ: Prentice Hall, 1986, 471 pp.
- Richard Behard: "Appendix A: Part Reviews: Disassemblers, Decoders." *Microprocessor Report*, December 9, 1992. A report on Segay's Accolade.
- Stefan B. Bratt: "Virtual Device Driver Support on Windows NT." Microsoft NT Device Driver Developer's Conference, October 1992.
- Douglas Boding: "Background Copying Without OS 2." *PC Magazine*, 17 January 1989. A DOS utility tasker, ISR which copies files in its background and discusses INT 28's and Oct. Sec. PSP's.
- Douglas Boding: "I'd Like to Propose: Give Yourself a Smart DOS Control Interface." *PC Magazine*, 26 December 1989. Discusses 28.5, Oct PSP, an DOS 2.x loading COM MANIFEST.COM's PSP, and other goodies set by OS, INT 11, status, etc. (DOS 4+).
- Douglas Boding: "Strategies and Techniques for Writing State-of-the-Art ISRs that Exploit MS-DOS 5." *Microsoft System Journal*, January/February 1992. What ISR writers need to know about state-of-the-art task switchers, and other new developments.
- Douglas Boding: "Quick Root Your PC and Protect Your Data with BOOTDIS and BOOTHPW." *PC Magazine*, March 30, 1993. These utilities use the DOS boot record.
- Paul Bourke: "Weird as Questions and Answers." *Windows DOS Developer Journal*. This excellent monthly column whose authors, on a Microsoft-sponsored open access basis of responses to DOS programmers, for example, the December 1992 column showed how to turn a Windows DOS box's system machine VM handle from its window handle. (HWSND).
- James W. Brock: "Structural Methodology: Technological Performance and Predatory Innovation: Relevant Standards under Section 7 of the Sherman Act." *American Business Law Journal*, Fall 1985. This work on Kodak's relations with its competitors has some interesting parallels to Microsoft's role in the software industry.

- Ed Brooks, *The Mind of Man Month: Essays on Software Engineering* Reading MA: Addison-Wesley, 1989, 268 pp. Many developers refer to this book, but few seem to have actually read it or paid attention to it. Brooks's writing style is good. Note Brooks's warning regarding David Parnas's then new concept of "formal or linking" as presupposed complete documentation. Chapter 15 discusses documentation, note the picture of Stonehenge: "The world's largest undocumented computer".
- See Brooks's also frequently Asked Questions list for newsgroup comp.os.msdos.programmers. Available at <http://www.cba.hawaii.edu/~pet/manager/m1/cdo/>. The DOS FAQ contains handy answers to some frequently asked questions. What the heck is IOCB? P. 646. "Wild Borland C code and Micro-Soft's code link together. How can I tell if a put or output has been redirected?," "How can I read or write in DOS's CMOS memory?," and even "Why won't my code work."
- Ralph Brown, ed. Brooks, ed. *Intercepts: A Programmer's Reference to BIOS, DOS, and Third Party Code* Reading, MA: Addison-Wesley, 1991. This massive book lists almost any imaginable procedure or routine for the PC, so long as it in one way or another works off an INT call. Even the BIOS's interrupt list is a delightful readable outline, and is included on the disk accompanying *Intercepts and DOS: An Interrupts and BIOS Reference Guide*. In 1994 Addison-Wesley will publish the second edition, and a new book, *Network Interrupts*.
- Dave Book, "DOS File Handle Limits," *11/11 Specialist*, February, 1993. This magazine is now *Windows, DOS Developer's Journal*.
- Chris, "Millions of easy MIBOS: Making getting system BIOS emulation code available by ftp from source code," DOS FAQ as the network redirection module.
- Cliff, *Cliff's Assembly Language Reference for the PC* CA: MISI Books, 1989, 337 pp. Assembly language source book for a network based on DOS device drivers.
- Geoff Chappell, "Specialty Section: INT 13," Apr. 1991. A critique of some subtle and not so subtle errors in the first edition of *Undocumented DOS*.
- Geoff Chappell, *DOS Internals*, 2nd Edition, MA: Addison-Wesley, 1993. Chappell's book is the natural follow-up to *Undocumented DOS*, covering many topics in more detail and more accurately. It will also help the classic case of variations of starting the system "footprint" extensive coverage of INT 19, the BIOS sector-by-sector interface, DOS in ROM, Windows/DOS interaction, BIOS interrupt services, state functions, writing ISRs and device drivers in C, XMS, or other chapter on the A20 gate, GADGET, a how-to look at HIMEM.SYS SmartDrive, how to look at some source code examples, BIOFILE, and an entire chapter on error codes. An amazing amount.
- Geoff Chappell, "Unraveling SMARTDrive," *The Windows Journal*, January, 1992. This article revealed the SMARTDrive BIOS sector interface to the SmartDrive 4.x "BAPI" interface (see the interrupt list on disk).
- Laura Chappell, *Networks and NetWare: A UNIX Analysis*, Selby, Newell Press, 1993. Includes a good look at the "signature" NetWare Core Protocol (NCP).
- Kyle, *Windows 3.11: The User's Manual*, Simon & Schuster, *Developing Applications Using DOS*, New York: NY: Wiley, 1990, 873 pp. This book by three IBM employees who were "the lead developers of the operating system" describes many undocumented DOS features, such as the "linking" of the DOS 4.0 system, "describes many undocumented DOS features, such as the 'linking' of the INT 21h function" 510.
- Peter, "The Unofficial DOS Programmer's Guide: Windows 3.1," *The Doble's Journal*, February, 1992. Includes a structured table of changes to the CallDword() pcx function.
- Robert Colson, "The GADGET Instruction," *11/11 Specialist*, October, 1991. This magazine is now *The Windows, DOS Developer's Journal*. Colson uses an ancient emulator (ICE) to take apart the undocumented GADGET instruction as a very good article.



- David Coates. "C/P/M 86 vs. MSDOS: A Technical Comparison." *The Hobby Journal* 11: 1982. An early look at the similarities and differences between C/P/M and MS-DOS 1.0.
- Helen Coster. *Inside Windows NT*. Redmond, WA: Microsoft Press, 1993. 385 pp. As the author puts it, "this book explains 'exactly' how NT 'sort of' works. Regardless of whether NT is appropriate for more than a small handful of users, the book nevertheless is essential reading for anyone interested in operating systems and how they work. The chapter on "protected subsystems" discusses virtual DOS machines.
- Ralph Davis. *Windows Network Programming: How to Survive in a World of Windows, DOS, and Networks*. Reading, MA: Addison-Wesley, 1993. 562 pp. An in-depth look at Windows for Workgroup, NetWare, LAN Manager, and other topics. Chapter 7 covers API1 a slat on a slat DPMI.
- Harve Deitel and Michael Kogua. *The Design of OS-2*. Reading, MA: Addison-Wesley, 1992. 389 pp. Chapter 10 on compatibility covers the OS-2 1.x and 2.x DOS box.
- DOS Protected Mode Interface (DPMI) Specification, Version 1.0*. March 12, 1991. Available from Intel under no. "40977-001". The 1.0 specification is more readable than the 0.9 version (document no. 240765-001), but be aware that 0.9 is what Windows implements.
- Ray Duncan, ed. *Advanced MS-DOS Programming*, second edition. Redmond, WA: Microsoft Press, 1988. 669 pp. This is the bible of DOS programming, with examples in assembly and C. Includes a discussion of how MS-DOS is loaded.
- Ray Duncan, ed. *The MS-DOS Family-pack*. Redmond, WA: Microsoft Press, 1988. 1,570 pp. Part C of this massive handbook, "Customizing MS-DOS," is particularly good, containing chapters on TSRs, exceptions, loaders, hardware interrupt handlers, DOS filters, and installed-device drivers. Includes Richard Vector's "dynamic piece on ISR programming. The article on "The Components of MS-DOS" describes the DOS boot sequence.
- Ray Duncan, et al. *Extending DOS*, second edition. Reading, MA: Addison-Wesley, 1992. 538 pp. Discusses DPMI, VCPI, VMS, IOPL, and 32-bit protected mode programming for DOS, DesqView, and other topics.
- Ray Duncan. "Microsoft Windows: 386 Creating a Virtual Machine Environment." *Micro Systems Inc. Conference* 1987. His comments an excellent description of how Windows 1 enhanced mode programs (multitasking, multiple DOS boxes) on top of a single copy of DOS. Discusses instancing of DOS data structures.
- Ray Duncan. "Programming Considerations for MS-DOS 5.0." *PC World* (October 29, 1991). Part 1, and November 12, 1991. Part 2. If you're checking the DOS version to test EMMS and the HIMEM.SYS, documented INI "Hardware task switcher" API.
- Martin, R.M. D. is an OS-2 to LAN Manager, in Stephen A. Kochan and Patrick H. Wood, eds., *LAN Networking and Networks*. IN: Hayden Books, 1989. pp. 237-284. You wouldn't know it from the title, but this discusses the Server Message Block (SMB) protocol including Extended SMB.
- Jeff Dantzman and Karl Weiskamp. *PC Languages: C, C++, Pascal, Basic, Java, Teacup, and Turbo Pascal*. New York, NY: Barron's, 1992. 626 pp. Includes section's aspects of site for ISR programming, calling DOS from Windows, moving DOS programs to Windows, working with files, code options, etc. in date, time manipulation, and backup and printing.
- Grant Feards. "OS-2 NetWare Workstation: Set INI 21h Functions." 1990. unpublished.
- Lam Eric. *Understanding VxDs: A Programmer's Guide to Revised NTWare, API and Protocol*. Reading, MA: Addison-Wesley, 1994. This book will cover the NetWare Core Protocol (NCP), the NetWare File API (NFI), modification of INI 21h, the FS interface, server APIs, and so on.



- IBM *Tools of Reference - Personal Computer AT* 1985 (IBM order number 6280070 and supplement 6280099). This is the IBM manual with the full BIOS source code listings. Apparently it's still available for about \$200 from 1-800-476-7287.
- IBM's OS - *IBM Central Program Programming Guide* (June 1992). This IBM documentation includes information on long features, extended attributes, EAS, the STAR, DATA structure, etc.
- IBM OS - *Virtual Device Driver Reference* IBM publication #s10g 6310. Part of the OS 2.2.0 Device Driver kit (DDK), there's also an OS 2.2.1 DDK.
- Intel *System Processor User's Manual* 1993. Appendix H "Advanced Features" refers to the *Supplement to the Platform Processor User's Manual* "available with the appropriate volume license agreements in place."
- David Lewis "Adapted Status of DOS" *Programme NOU-UK* April 1993. A utility in C for editing the master environment.
- Steve Lewis "DOS Meets Real Time" *Embedded System Programming* February 1992. How to build a real-time DOS. See also Steve's manuals for General Software's Embedded DOS products.
- Richard Kitson "Command Line Tricks of MS-DOS" *Programme NOU-UK* June 1991. Using the 2E AF interface to add new "internal" commands to COMMAND.COM.
- Rick Koldzschig "Locate Available IRQs with FINDIRQ" *PC Magazine* September 28, 1993. Describes IRQ programming, and a practical use for waking the DOS device driver.
- Thomas Kouri *Let's Get Programming* Stamford, CA: Center for the Study of Language and Information, 1992. 468 pp. Kouri treats programs as a new form of literature. Except for programming, also requires a form of literary criticism. In a odd way - by taking apart MS-DOS and Windows, we are solving the world's most prevalent operating systems to see for real "every of them."
- Kim Kokkonen's "SR utilities" 1989. Kokkonen provides the Turbo Pascal source code for the following executables: SR utilities, MARK, RELEASE, MARK, MARKNET, RELEASE, WATTEH, DISM21, RAMREF, MAPMEM, DDIR1, EMMEM. Available from Compuserve, or the BPROC.A forum.
- Thomas Krauthammer and Steve Szepor "An Incomplete Education: Raising Real Costs to Achieve Power over Power" *Yale Law Journal* December 1990. Read this - each as a complete treatise on what Microsoft is trying to achieve with its core stress delivery of its reusable executables. Or perhaps it's at least Schwarzenegger's writings - a "caveat emptor" instead.
- Jim Kye "Application Wappers" *PC Techniques* June/July 1992. Writing applications that spawn other applications as an alternative to writing OSs.
- Robert S. Leland *The Windows Group: Behind MS-DOS (2nd Edition)* second edition Reading, MA: Addison-Wesley, 1992. 560 pp. The second Edition is substantially the same as the first (1987) except for the addition of a chapter on CD-ROM device drivers and writing DOS device drivers in C. It is a standard reference on writing DOS device drivers.
- Gene K. Lerner *In Source: The Code and Methods of an IT Corporation* Reading, MA: Addison-Wesley, 1993. 548 pp. This book/disk package (the disk contains sample forms, queries and agreements) is superb. Includes words of trade secrets, reverse engineering, the use of source code reviews, strategies and everything you as software developer would want to know about the law.
- Murray J. Lester "Extending COMMAND.COM" *Windows DOS Developer Journal* February 1993. The 2E AF installable command interface.
- Gregory C. Little *Inside OS-2* Redmond, WA: Microsoft Press, 1988. Discusses backwards compatibility, de facto standards, etc.
- David Long "ISR Support in Microsoft Windows Version 3.1" Microsoft Developer Network MSDN, CD-ROM (see below).

- Steve Marz and Paul Andrews.  *Gates: How Microsoft's World Reshaped an Industry—and Made History in the Rise of Man in America*. New York: Doubleday, 1993. 534 pp. This is not really the same as Gates's in-depth, 300+ page book of Microsoft's "how that interesting" Meticulously researched, with every fact and quotation backed up by at least one footnote. This book covers every important Microsoft purchase of QDOS, to its OEM pricing of DRPs, to how Murray Sargent in "Data Wars" moved Windows to protected mode.
- Steve T. Morison.  *Windows OS 2.0 Disk 1 Drivers in C*. New York: Van Nostrand Reinhold, 1992. 400 pp. C++ price of obscure OS 2.03 (aka "sketchy drivers" VDIS) and the Virtual DOS Machine VDM. Also, selection for OS 2.2, should be available.
- Michael Murray. "The PE File Format on Topview: sort of Emacs." *Dr. Dob's Journal*, July 1993. Long and informative. Lists PE's and what Windows uses to run "old" DOS programs.
- Michael Murphy. "Choosing CONIFER SYN Options at Boot." *PC Magazine*, November 29, 1988. Discusses the undocumented CONIFER SYN buffer.
- Michael Murphy. "Running Programs Fearlessly." *PC Magazine*, February 16, 1988. Discusses the problems with using INT 2Fh.
- Philippe Mercier. *La maîtrise des programmes exécutés sous MS-DOS*. Marabout, 1990. 410 pp. A handy book on ISRs from Belgium. Did you know that the French for "botkey" is "touche magique"?
- Microprocessor Report.  *Understanding the 80486 Microprocessors*. Emeryville, CA: Ziff Davis Press, 1993. A 500-page book on the 386, 386, 486, and Pentium from Michael Slater's brilliant newsletter. *Microprocessor Report: The Computer's Guide to Microprocessor Hardware*. Includes discussions of cache, multitask processor, instructions, a solid online section on legal issues. Total 500+ articles. Also covers the old articles by John Wharton, such as "Commod Marketing." Why can't all technical writing be like this?
- Microsoft. API to Interact with MS-DOS Instance Data. Undated internal document. Describes the DOSMGR call or API.
- Microsoft. *Device Driver Adaptation Guide*. DDMG, version 3.1, Redmond, WA, 1992. Included with the Windows 3.1 Device Driver kit. DDMG also includes an essential appendix on "Windows Interrupt 2Fh Services and Notifications." When they've put this important stuff in an obscure manual like this is because... There's also a useful chapter on Windows network drivers.
- Microsoft Developer Network. MS-DOS CD ROM. A must have for any serious DOS or Windows developer. The MS-DOS CD ROM includes huge amounts of information that programmers often mistake for junk, is undocumented. For more information, call 800-759-5474 ext. 206-936-8661. There's a very nice sample of the articles related to DOS. Determining Windows Version Mode from MS-DOS Apps. Disabling Paging MS-DOS Applications, "Global ISR Pop-ups Incompatible with Windows 3.11 Screen DOS Apps Show Error Messages in Enhanced Mode." Do Not Use the MS-DOS API Set Utility in Windows, "Calling a DLL Written for Windows from a ISR for MS-DOS." Finding a ISR to a VxD, "Using the Interrupt 2Fh Critical Section Semaphores." How a ISR Can Semaphores Access to Its Data. "IO of Calls in Protected Mode Microsoft Windows." Access to the Windows Clipboard by DOS Applications. "Windows 3.1 Standard Mode and the ACPI." How Microsoft Windows Uses an MS-DOS Mouse Driver. How to Start a Windows Application in Directx from DOS, "Passing File Information a ISR to a Windows Application." This partial list of titles should make clear that MS-DOS documents are so vast and varied, all and even the hard DOS programmers can't ignore Windows.
- Microsoft. Statement on the Subject of Undocumented APIs. August 31, 1992. Microsoft's news release on the "secreted Windows." There are undocumented APIs in every major operating system and applications developers routinely make use of them." At the same time, Microsoft issued

"Questions and Answers About Documented and Undocumented APIs," and a ready-to-use white paper entitled "Undocumented Functions."

- Microsoft, "MS-DOS API Extensions for DPMI Hosts" version pre-release 002, October 31, 1990. A Microsoft confidential document.
- Microsoft, *MS-DOS Programming Reference*. Redmond, WA: Microsoft Press, 1991. 464 pp. The official programming reference for DOS 5. Microsoft advertised this work, which came out shortly after the first edition of *Undocumented DOS*, as "DOS Documented." For the first time, Microsoft officially documented some well-known undocumented functions, while still leaving many holes. While an essential reference, the book has some omissions, which are explained in chapter 1.
- Microsoft, *MS-DOS Programming Reference*. Redmond, WA: Microsoft Press, second edition, 1993. 512 pp. This is the MS-DOS6.0 version of the programming reference. It corrects some of the errors from the DOS 5.0 version, and adds chapters on DoubleSpace and MRU.
- Microsoft, OEM Adaptation Kit, OAK. Personal Computing Update for MS-DOS 5. November, 1991. Published through the OAK provides a wealth of information on DOS interfaces, including discussions of DOS 5, virtualization, BIOS support, OEM customization of DOS, a MBR, ROM boot loaders, DOS in ROM, 21-bit of APM and how to use the PMMCA, a source code for disk loaders, SNA, AR, ASM, SNA, VAX, structure, WPA, CHTRN, DOS patches, DOS IAR, ASM, more DOS patches, PDB, ASM, the PSP structure, CONN12, ASM, DOS, vga, MUI, ASM, INI, 21-bit calls, including the network reference, and OEMNUM, INC, OEM number assignments.
- Microsoft, *Virtual Device Adaptation Guide*, VDMG, version 3.1. Redmond, WA, 1992. Included with the Windows 3.1 Device Driver Kit (DDK). This is the definitive guide to VxDs and VMM. The DDK also includes source code for some of the VxDs that run in Windows.
- Microsoft, *Windows 3.11 Reference and How to Substantive Device Driver Code*. Redmond, WA, 1993. Included with the Windows NT DDK. It includes "Virtual Device Drivers for MS-DOS Applications that Use Special Hardware."
- Microsoft, *Windows Workbench, Release 3.11*. Redmond, WA, 1992. Writes says Microsoft "Secretly know how to produce good documentation. Its "reference kits" are excellent and often more detailed information that what the programmer's documentation.
- Microsoft, *Windows Release 3.11*. Redmond, WA, 1992. 538 pp.
- Microsoft, *Windows Software Development Kit*, SDK. Redmond, WA, 1992. The Windows 3.11 version of the SDK is generally quite good and describes many formerly undocumented interfaces. The *Programming Reference, Volume 1: Overview*, contains several chapters that are especially good for programmers, including chapter on network applications, and an extensive appendix and sample. "The book is organized and the portions are functional" chapter on "Windows Applications with MS-DOS Functions." Volume 4, *Reference*, describes the New Executive (NE) in format.
- Microsoft, "Windows 386 Paging Input Specifications" formerly Global FMM, 1990 spec, document revision 3.10.003, interface version 1.11, August 3, 1991. A Microsoft confidential document.
- Icd Mircea, "DOS Memory Control." *Tech Journal*, October 1987, p. 45. A good discussion of the layout of MEMB from the popular "Tech Notebook" as a new default manager.
- Icd Mircea, "Emulation 3214 of DOS." *PC Tech Journal*, February, 1989, pp. 229-233. Describes the structure of the DPMI.
- Icd Mircea, "More 11 index for New Applications" and "More Handles for Old Applications." *PC Tech Journal*, April 1988, pp. 161-165. Describes the file handle table within a process's PSP.
- Charles Miller, "Integrating a Windows Program to a Real Mode Device Driver." *Windows Developer's Journal*, August 1993.

- Paul Niles, *Network Programming in C*. Que, 1990. Includes a chapter on "Novell's Extended DOS Services."
- Paul Niles, "Self-Loading Device Drivers for DOS." *Windows DOS Developer's Journal*, May 1993, pp. 7-43. Another approach to loading device drivers from the command line.
- Robert F. Nims, *The Fundamentals of Computer Technology*. Boston, MA: Warren, Gorham & Lamont, 1988. See also 1991 Cumulative Supplement No. 1. Covers reverse engineering, antiroot, "site specific" virus eradication, "bring arrangements" documentation, obligations, mass market contracts, warranties, and more.
- David Newton, *Building Windows Device Drivers*. Reading, MA: Addison-Wesley, 1992. 434 pp. If you just want a general overview of the services that the Windows Virtual Machine Manager (VMM) provides (not device-specific VxDs) and don't want to buy the Device Driver Kit (DDK), this is MA (an inexpensive alternative).
- Nick O'Leary, "The Power of the NetWare DOS Requester." February 1993. Explains the limitations of the NetWare 2.11 network and describes the new DOS redirector-style "requester" in NetWare 4.x.
- Nolan, DR. *DOS 6.0 Optimization and Customization Tips*. September 1991 (earlier published by Dr. O'Leary). Dr. O'Leary's DR DOS 6.0 Version Numbers.
- Nolan, DR. *DOS 6.0 and the Internet*. *Guide*. DR product #1182. 2013-002.
- Nolan, Dr. *NetWare and the Internet*. *Overview*.
- Tommy Olin, "Making Windows and DOS Programs Talk." *Windows DOS Developer's Journal*, May 1992. A series of real approaches to DOS/Windows communication, including the 21-17 chip board API and a pipe interface VxD.
- Walt Olin, "Communicating Between Virtual Machines." *Win/Dos East 1993*. April 26-30, 1993, Boston, MA. Boston University Corporate Education Center. Another great presentation by Walt Olin, formerly of Rational Systems.
- Walt Olin, "Using the 21-17R." *Windows Magazine*, November 1991. How to use 21-1605 to minimize disk I/O to properly maintain costs/statistics/Windows.
- Walter Olin, "Programming for DPMI compatibility." *Software Development '91*.
- Walter Olin, "Using DPMI to Hook Interrupts in Windows 3." *The Dobb's Journal*, February 1992.
- Orlando Sykes and Wally, "Regulatory Systems Review: A Reply." *Columbia Law Review*, June 1983.
- Earl P. Peterson, *Archives from CMS/DOS*. *Box*. June 1983.
- Frederick P. "The MS-DOS Debugger Interface." See Schindler et al. *Undocumented DOS*, first edition. Reading, MA: Addison-Wesley, 1990. This chapter was dropped from the second edition, because Microsoft has documented 21-4304. Load But Don't Execute. But the original documentation is so scary and detailed the DOS 6 programmers references is wrong that you'll still use it. This chapter *should* be the first edition. I can warn everybody anything with 21-4304 which is essential to DOS developers. I. DOS 5.25's debugger, a debugger would also need to be aware of 21-4304. See Executive Summary, see Cripps's *DOS Internals*.
- Charles "Good" Wheeler, "The Path of Magic." *Mag*, no. 28 April 1998. Discusses "the undocumented and strange" INI 21-b.
- Philip J. Software, *MS-DOS 5.0 and 6.0: Developer's Guide*. The writing in this manual is curiously similar to parts of *Undocumented DOS* and *Undocumented Windows* and contains protected mode versions of several programs (but no undocumented DOS).
- Mark Sikes, *Windows Internals: An In-Depth Look at the Windows Operating Environment*, Reading, MA: Addison-Wesley, 1993. In chapter 1, Mark shows how Windows boots on top of DOS, with a particularly detailed look at WIN.COM and the Windows KERNEL.
- Scott Pak, "Reverse Engineering Reverse." *Topik*, May 1993. A report on the Sega, Accolade and Nintendo's Atari cases.
- Mike Pastorefsky, *Disrupting DOS: A Code Level Look at the DOS Operating System*, Reading, MA: Addison-Wesley, 1994. Mike is a author of a DOS workalike called RxDOS. This book presents

- the assembly language source code for BIOS providing a unique inside look at how a BIOS really works
- Jeff Prowse. *IO's & Techniques and Utilities*. Emeryville, CA: Ziff Davis Press, 1993. 1035 pp. This massive book for both users and programmers includes chapters on writing ISRs and device drivers.
- Jeff Prowse. "Hidden Goodies in DOS 5.0." *PC Magazine*, April 27, 1993. Brief explanation of 21 4A HMA functions, of EMM386 EXECPI support, and of the FDISK, MBR switch.
- Jeff Prowse. "How Device Drivers Work." *PC Magazine*, November 28, 1989. Finding the device chain via 21/52.
- Jeff Prowse. "The Inner Life of a ISR." *PC Magazine*, August 1990. Discusses the InDOS flag, INT 28h, critical error flag, and Get/Set PSP.
- Jeff Prowse. "Instant Access to Directories." *PC Magazine*, 14 April 1987. Excellent discussion of ISR programming, discusses the InDOS flag, INT 28h (including the need, not only to hook INT 28h, but also to periodically invoke INT 28h as well), the DOS stacks, I/OA critical errors, and hooking the BIOS disk interrupt.
- Jeff Prowse. MS-DOS Q&A column in *Micro Systems Journal*. Sample columns: May/June 1992 on MIB chain and first I-MIB; M. signature; Sept 1992 on SMART Drive 4 interface; Dec 1992 on making a program read itself high; and on detecting RAM errors; March 1993 on 21 4A HMA functions; May 1993 on 21 4103 load works, critical errors, and DoubleSpace; June 1993 on 8514 A detection, Workvace, and the critical error flag.
- Jeff Prowse. *PC Magazine IO's & Utilities Management with Utilities*. Emeryville, CA: Ziff Davis Press, 1993. 405 pp. Includes assembly language source for UMHIIFIS, FILEMON, HMACALL, INSIAC, RIMOV1, and other utilities.
- Jeff Prowse. "Replicating Internal DOS Commands." *PC Magazine*, December 31, 1991. Using the 2E/AF interface.
- Jeff Prowse. "Teaching a ISR New Tricks." *PC Magazine*, 12 June 1990. A nice discussion of the active PSP.
- Jeff Prowse. "What HIFIS Does." *PC Magazine*, November 12, 1991.
- Jeff Prowse. "Unlocking the Mysteries of CHIKDSK." *PC Magazine*, May 11, 1993.
- Quantum Corporation. *Using ISR Tools*. Report 200A. A VSR toolkit for assembly language and C programmers.
- Robin Raske, Charles Petroz, and Stephen Rayne Davis. "Taking Up Residence." *PC Magazine*, November 25, 1986. Singing the "ISR blues."
- Frederick Raymond, ed. *New Hacker's Dictionary*. Cambridge, MA: MIT Press, 1991. Some will find this book annoying, but it does have good descriptions of several important pieces of program interface, such as "You're not expected to understand this" and "I am booksy."
- Tom Rizzo. MS-DOS CD-ROM Extensions. A Set-In-Place Access Method. *Micro Systems Journal*, September 1988, pp. 54-60. Describes the Microsoft CD-ROM Extensions (MS-DEX), including its implementation of using the DOS network redirector.
- Jeff Roberts. "10 Steps to Windows Coexistence." *PC Techniques*, February/March 1991. "Take your DOS by the back of the neck and shove it into Windows."
- Wendy Cordina-Robey. "With the FCC, Come to its Sense About Microsoft's Mischiefs." *Uprobe*, August 1993, pp. 11-27. A scathing piece on Microsoft and the FCC, by the author of a forthcoming book on the subject. One *Undocumented DOS* contributor appears in this article under the name "Michael Hatch." Includes a telephone interview with Brad Silverberg, Microsoft VP of Systems Software, regarding the VARD code, see chapter 1. Also see the letters to the editor at *Uprobe*, October 1993, which were overwhelmingly pro-Microsoft: "In a free market, Microsoft can design its products to do anything it pleases." "What right does the FCC have to regulate Microsoft?" etc.

- Charles Ross, *Programmer's Guide to NetWare*, McGraw-Hill, 1990
- Nancy C. Ross, *Principle of Ancient Law*, Westbury, NY: Foundation Press, 1993, 606 pp. An excellent, up-to-date, and post Reagan survey of the field including coverage of "Market Structure and Monopoly Power," Berkeley v. Kodak, "predatory innovation" (Transamerica v. Fireman's Fund), integration, exclusionary conduct, vertical restraints, tying arrangements, and so on.
- Victor Rodin, "Walking the OS 2 Device Chain," *Dr. Dobbs' Journal*, October 1990. An interesting article on walking the device chain under MS-DOS.
- Scott Rosenberg, *PC: Mastering Turbo Pascal with Techniques and Utilities*, Emeryville, CA: Ziff Davis, 1991. Objects for accessing undocumented DOS.
- Scott Rosenberg, *PC: Mastering Turbo Pascal to Windows Techniques and Utilities*, Emeryville, CA: Ziff Davis, 1992, 1,100 pp. Calling undocumented DOS from protected mode IPW programs, includes a chapter on "Access to Real Mode."
- Richard Sadowsky, "Attacking DOS Fragmentation," *Windows Tech Journal*, April 1992. Accessing a selected real DOS from a selected mode Windows program.
- Richard Sadowsky, "Is a Real Mode Out There," *Windows Tech Journal*, January 1992. From the editor's perspective, but more relevant than *Word & Tech Journal* articles are about. This one's about using DPMI to real mode, get it. IFR communications with Windows.
- Richard Sadowsky, "A Primer on Protected Mode," *Windows Tech Journal*, January 1993. Includes a correction to Sadowsky's January 1992 article.
- Lee Saxe, "An Exception Handler for Windows 3.1," *Dr. Dobbs' Journal*, September 1992. Assembly language source code for a VxD, WIN32.386.
- Murray Sargent and Richard Shonemaker, *The PC: From the Inside Out*, Reading, MA: Addison-Wesley, 1994. This is the third edition of Sargent and Shonemaker's famous *The IBM PC: From the Inside Out*, the new edition will cover protected mode, the post IBM PC architecture, 386, 486, Pentium machines, Windows VxDs, and a host of other topics.
- Arno Schaefer, *IBM's VxDs: Programmer's, The Industrial Reference*, Boston: Addison-Wesley (Osmanov), 1991, 112 pp. This massive book includes a lot of material on DR-DOS.
- Roger L. Schuchart, *Legal Aspects Patents and Intellectual Property*, St. Paul, MN: West Publishing Co., 1986, 272 pp. A quick and dirty but covers competition, trade secrets, the FTC, warranties, and other topics.
- Robert Schuchart, "Supercharging IFRs," *Books & Code in C*, Berkeley, CA: Osborne/McGraw-Hill, 1989. Chapter 3 presents a useful IFR skeleton in Turbo C.
- Andrew Schuman, "Accessing the Windows API from the DOS Box," *PC Magazine*, August 1992. Part 1, September 15, 1992. Part 2, and September 29, 1992. Part 3, "Using the Windows 2B/17 clipboard API from a DOS program."
- Andrew Schuman, "The Programming Challenge of Windows Protected Mode," *PC Magazine*, June 25, 1991. Accessing undocumented DOS from a protected mode Windows program, using DPMI and source code for undocumented Windows APIs calls.
- Andrew Schuman, "Can I do Anything and Do Anything with 32-bit Virtual Device Drivers for Windows," *Microsoft System Journal*, October 1992. Describes VxD APIs accessible via 2B-1684.
- Andrew Schuman, "Expanding a Standard Paged Virtual Memory in Windows Enhanced Mode," *Microsoft System Journal*, December 1993. Discusses DPMI and VxDs.
- Andrew Schuman, "Call VxD Functions and VMM Services Easily Using Our Generic VxD," *Microsoft System Journal*, February 1993. This article describes in more detail the generic VxD used in *Undocumented DOS*, chapter 3.
- Andrew Schuman, "Walking the VxD Chain in Windows and Chicago," *Dr. Dobbs' Journal*, December 1993.
- Andrew Schuman, David Mavey, and Matt Pietrek, *Undocumented Windows*, Reading, MA: Addison-Wesley, 1992, 715 pp. By showing that Windows is just a fancy-looking DOS extender, in many



- ways seem just a new version of DOS, *Undocumented Windows* should give even the most die-hard DOS fanatic "I don't do Windows" a reason to look at Windows.
- Larry Schick, "Tough File Names are Not FAT Issues: Hope for Future DOS," *PC Week* September 20, 1993. An advance peek at the implementation of long filenames in Chicago.
- Mike Sorens, "The Undocumented LAN Manager and Named Pipe APIs for DOS and Windows," *The Dosh Journal*, April 1993. A survey of the 21.5k-calls-supported LAN Manager, the named pipe calls are supported by other networks including NetWare.
- Barry Simon, "Providing Program Access to the Real DOS Environment," *PC Magazine*, 28 November 1989, pp. 309-314. Discusses the three common DOS environments: the real or native environment, the program environment, and the real user environment.
- Bob Miska, Eric Straub, and Richard Freedman, *Inside MS-DOS 6.0*, July 1993. Three Microsoft developers provide a good overview of DoubleSpace internal data structures, although they fail to mention Vertosoft's who developed DoubleSpace in the first place.
- South Mountain Software, "Student C User Manual," Orange, NJ, 1991. 111 pp. A student manual for South Mountain's ISR Library, with a good explanation of ISR "theory." Also see South Mountain's ISRs and Hold Everything libraries.
- MSB Electronics, *Stacker 2.0 User Guide*, Appendix B "Programming with Stacker." This information is included in later versions of the Stacker user's guide; a description of a Stacker API is available on CompuServe (UCD5FAC).
- Al Stevens, *Inside C Memory Resident Utilities, Session 1.0 and Programming Techniques*, Portland OR: MIS Press, 1987. 315 pp. Chapter 11 "Memory Resident Programs" and Chapter 12 "Blinking Turbo C Memory Resident Programs" present a ISR skeleton. The author's later *Extending Turbo C Professional Port and OR: MIS Press*, 1989. 418 pp., contains an improved ISR skeleton plus a fascinating ISR user's manager.
- John Switzer, "Using DOS's Backdoor," *The Dosh Journal*, October 1990. Explains a bizarre quirk of the DOS Device I/O function: 21.13, with details on the INT 30h and INT 31h alternate function dispatchers.
- Panel Society, "The NetWare Core Protocol (NCP)," *The Dosh Journal*, November 1993.
- Andrew Tanenbaum, *Modern Operating Systems*, Englewood Cliffs, NJ: Prentice-Hall, 1993. Another great book by Tanenbaum.
- Andrew Tanenbaum, *Operating Systems: Design and Implementation*, Englewood Cliffs, NJ: Prentice-Hall, 1987. 719 pp. A superb operating systems textbook, with special focus on the important topics: the file system, and how on scheduling algorithms. Includes C source code for MINIX.
- David Thayer, "Less Filling, Tastes Great: A Programmer's Survey of MS-DOS Version 5," *Microsoft System Journal*, July 1991. See correction in *MSJ* November/December 1991, p. 1-6. Includes a discussion of IBM's proposed protocol for interrupt chaining.
- David Thayer, "Safe Functions for Working with MS-DOS Files," *Microsoft Systems Journal*, July-August 1992.
- David Tucker and Bryan Woodruff, *Writing Windows: Virtual Device Drivers*, Reading, MA: Addison-Wesley, 1993. The accompanying disk includes Microsoft's VxD files. See also Tucker's series on VxDs in *Windows Tech Journal*, March 1992-August 1992.
- Jean Van de, *The Theory of Industrial Organization*, Cambridge, MA: MIT Press, 1992. 479 pp. Who ever points Microsoft's role in the software industry is actually asking a question about "industrial organization." The math in this book is very difficult, but the material on tying arrangement, entry, network externalities, lock-in, product differentiation, incompatible standards, and bottlenecks is extremely useful.
- Paula Tomlinson, "The VxD Backdoor," *Windows DOS Developer Journal*, May 1993. As part of an article on NE device drivers, Tomlinson describes ring transitions via illegal instructions

80386. See the letters to the editor in the July 1993 issue pointing up the similarities to IBM's PASCADSP instruction thirty years earlier in System 360 and later in VM 360.
9. *How to Ask a Computer Vendor: PCs on the 386*. *Byte*, IBM special edition, Fall 1990. A detailed explanation of V86 mode, virtual machines, and device virtualization.
10. *Code Tester: Finding MS-DOS Device Drivers from the Command Line*. *ENR*, August 1989, pp. 16-20. TurboPower Software, *ENR and How*, Colorado Springs, CO, 1997, 353 pp. A detailed manual on how to use ENR in C++ using TurboPower's library, which comes with complete source code. ENR Made Easy and Object Professional for Pascal provide similar abstractions for Turbo Pascal programmers. Some of the examples are very involved. If you're writing ENRs, you really ought to be using *How to Ask a Computer Vendor* abstractions from South Mountain or Quantism.
11. *Using C to Write Multiple Virtual-DOS Machines: A Better DOS*. *an Disk Confusion*, ed., *OS/2 Note and Read*, Redwood Valley, Microsoft Press, 1990. Contrasts the OS/2 1.x and 2.x BIOS boxes.
12. *Communications Windows Source: Translating the Processor*. San Jose, CA, 1993. Manual by Andrew Scott, Ian, and Frank van Collier, includes in-depth discussions of "Understanding VxDs," "Understanding OS/2 Device Drivers," discussions of Windows executable file formats and seven
13. *How to Use C++ to Deconstruct and Fix*. Reading, MA: Addison-Wesley, 1993. This book is based on Frank van Collier's years experience of deassembling BIOSes. He is the author of the Source4 disassembler, which was included in product for deassembling BIOSes. Covers BIOS functions, absolute addresses, locations, I/O ports, and the CPU. Includes chapters on CPU detection methods, adapter card development, the BIOS database, the hard disk, the keyboard, CMOS, system time, DMA and I/O MCH, the interrupt controller and NMI, and system timers.
14. *Write a Super MS-DOS Developer's Guide*, second edition. Indianapolis, IN: Howard W. Sams & Co., 1989, 283 pp. Chapter 5 "Program and Memory Management," chapter 4 "Terminate the MS-DOS Resident Programming," chapter 6 "Installable Device Drivers," and chapter 11 "Disk Layout and File Recovery" all discuss undocumented DOS.
15. *Write a Super MS-DOS Paper*. Indianapolis, IN: Howard W. Sams & Co., 1988, 578 pp. Some chapters have several chapters, also excellent discussions of undocumented DOS, chapter 6 "Random J. Mikes' Undocumented MS-DOS Functions," chapter 7 "M. Steven Baker 'Safe Memory Resident Programming' ISR," and chapter 10 "Walter Dixon 'Developing MS-DOS Device Drivers'".
16. *Roger Ward, ed., MS-DOS System Programming*, 2nd Edition. Lawrence, KS: R&D Publications, 1990, 239 pp.
17. *Williams' DOS and Windows Protected Mode*. Reading, MA: Addison-Wesley, 1993, 543 pp. An overview of protected mode DOS extenders for the purposes of this book. Windows is just another protected mode DOS extender. Includes DPMI programming, interfacing ISRs to Windows 32 bit programming.
18. *Al Williams' DOS's Low-Level Guide*. Redwood City, CA: M&L Books, 1991, 914 pp. Includes sample code on ISR programming, "Big Real Mode," and a protected-mode DOS extender.
19. *Richard Wilton, "Create and Use Resident Utilities"* in Ray Duncan, ed., *MS-DOS Fun!lopedia*, pp. 147-185. This article carries the note that "Microsoft cannot guarantee that the information in this article will be valid for future versions of MS-DOS." Nonetheless, it is probably the definitive article on ISR programming, interestingly, Microsoft eventually documented all the then-undocumented functions discussed in this article.
20. *K. L. Zerkov, "Fix Windows Virtual Machine Control Block"*, *Dr. Dobbs' Journal*, forthcoming in January 1994 (Part 1) and February 1994 (Part 2).

- 26, n48  
 Add instance from, 42  
 aim blocks, n3, n5  
 base, requires, 59  
 BCDL ASCII, n3  
 chain attr. > 27, 90, 513, 553  
 check, 558  
 dos allocation, n2  
 dos findfile, n2  
 dos getdrive, n5  
 dos getvch, n7, n2  
 dos keep, 553, 54  
 dos loadfile, n31  
 dos open, 486  
 dos setdrive, n2, n5  
 dos server, n7, 544, 552, 579  
 format, n39  
 Get EMS VMS Limits, 43  
 get var handle, 156  
 GetL (off) attr, 44  
 MM VAR, n3  
 name, n6  
 name, n396  
 pop, 286  
 SelectionMapHit, 125  
 set stack, 556  
 NetMapping Info, 63  
 Test Global Mem, 42
- X**
- X20 state, 289, 90, 294, 296, 309, 349, 50, 438  
 XARD detection code, 7, 18, 185, 191, 489  
 test for I/O legitimacy, 40, 11  
 utility, dos objv, and, n7  
 XEN CCM and, 8, 10, 14, 16
- Addresses**  
 base, 112  
 base, 112, 137, 155  
 allow, real seg., 149, 51  
 Address Pbl C in Block, 123  
 AlsoSelector, 130, 131, 137, 138, 142  
 ALWAYS Field, 35  
 Andrews, Paul, 46, 161  
 Ann Varus, ALWAYS, 39, 43, 45  
 Antitrust, 10, 3  
 API Application Programmer's Interface, 101, 179, 609  
 C, e.g., dos initialization and, 424  
 C, coding, 195  
 C, P, and, 181
- DISQView and, 17  
 Accessing, instructions and, 50  
 DOS emulation and, 209  
 DOSM, R and, 74, 10, 124, 126  
 DoubleSpace and, 481, 485  
 I/O and, 107, 170, 40, 41, 141  
 I/O times and, 112  
 I/O and, 78  
 N, 170, and, 197, 204  
 198, 2 and, 214, 214, 16  
 ports and, 103  
 M, W, K, C, and, 469  
 Stack, and, 460  
 User-data service, 124  
 users and, 149  
 Windows NT and, 218, 219, 220, 221, 225, 226  
 APPEND, 446, 467, n30, n55  
 application, 548, 607  
 Application wrappers, 95, 99  
 ARMA, 49  
 ARM code, 186, 191, 222  
 command interpreters and, 640  
 646, n51  
 resource management and, 345, 351, 353  
 Assembly language, n1, n2, n5, n8, 11  
 See also Disassembly  
 compiler rules for, 64  
 I/Os and, 551, 62  
 codes and, I/O calls from, 77, 74  
 VMS, N, 446, 467, 546, 548  
 VMS, C CMM, 446  
 atan, 377  
 AWK, 37, 24
- B**
- Back end redirectors, 497, 98  
 BARCHI, 406, n  
 BASIC, 182, 60  
 DOS calls from, n6, n7  
 undeclared I/O calls from, n4, n5  
 Begin Critical Section, 37  
 Begin Next VMS Exec, 125  
 BeginPanic, 103  
 Beta testing, 7, 10, 15, 16  
 Beta V, 451  
 BkAdvs, 20, 21  
 Boring, Douglas, 50  
 Bonczak, Paul, 166, 169  
 Boot records, 408, 11  
 BootSwap, 17
- BPB, BPD Parameter Blocks, 410  
 11, 413, 449, 50, 498  
 BRL AN, C, 563, n4  
 Break information, 563, n4  
 Brown, Bob, 44, n9, 173, 73  
 Roll Back + Interrupt Log, 44, n9  
 Buffers, 42, 87, 406, 436, 43  
 BUIBRS setting, 42  
 BUIBRS+ 406, 437, 438, 4  
 BUIBRS+ CMM, 87  
 build, n5, n8  
 BuildVIB, 38
- C**
- C, C++ programming, R, 122, 173, 195, 179, 542 See also Quick  
 VARD detection code and, n6  
 application wrappers and, 97  
 disassembly and, 69  
 I/O calls from, n2, n5  
 DOS emulation and, 209  
 I/O and, 50, n1, n7, n3  
 object functions and, n5  
 printf, and, 104, 186  
 QWINS and, 47  
 truncates and, 148  
 VMS programming in, 548, 550, 56  
 undeclared I/O calls from, C, 76, 8  
 warnings and, 199, 206
- C, C++  
 DOS systems and, 402, 436, 43  
 C Ad, 642  
 C All ASSIGN, 11, n7  
 C All IN, 11, BRL P1, n7, 35  
 calls, 365  
 cd, 522, 23  
 C D BMM, 40, 45, 40, 192, 593 See  
 also MM DIX, C D BMM versions  
 accelerators, 518  
 ANS, N and, 446  
 C V: file format and, 451  
 drive letters and, 462  
 logical drives and, 69  
 redirector overrides and, 494  
 speedup products, 441  
 tracing an open and, 503  
 C DS of server Directories Structure, 2  
 10, 42, 93, 101, 2, 185, 498  
 attr, with, the, 448, 49  
 C DS IIN, and, 22, 73, 74  
 C DS Size and, 30  
 changing I/O and, 87, 88







generate addresses 512  
 GetNtRtlSyntax, 245  
 getnterraps 380  
 Get/Set System, 287 49  
 Get Cur Vm Handle 22 23 156  
 get dphs : 415, 416, 420, 430, 434  
 450, 464  
 get fat entry 416, 420, 430  
 get index, 121  
 get mapped 137, 140  
 get mch, : 354  
 Get Next Vm Handle 156, 157 160  
 get pba, 433  
 Get Real 481  
 get real address, 141 148, 481  
 get sda : 108 121 122  
 get section 519  
 Get Sys Vm Handle 156, 157  
 get system, 106 122  
 GetAppOutputFlags : 75  
 getbks : 65  
 Get First Mch 351  
 Get Out 485  
 getps 640, 658  
 GetFileStat, 87 GetMessage : 103  
 GetNextMHandle 159 160 162  
 GetNTLDR 108 9 109  
 GetNtControlBox 131  
 GetNtMHandle 159 160  
 GetLastQueue 47  
 GetVADP1 162  
 GetVersion 166  
 GetWindowsTask 166  
 GetWinSys 38  
 GlobalDPCs, : 37 130  
 GlobalDPCs, : 130  
 GetPools 110 13  
 GetPWA k and 143  
 Windows DPCs extenders and 121  
 GRAPHC 5, 546  
 GrootStToMax 37 477 482

## H

h2name 474  
 H2NtAPI 472 74 476, 483 488  
 H2NtAPI : 472 74  
 Hae John, 209  
 HAVe2E ASm, ntd  
 HPI31C 231 34  
 "Hello world" 103  
 HDOS, 191  
 HDOS SYS, 182 190  
 HDMA, 190  
 HMEM SYS, 7 8, 10  
 HMA, High Memory Areas, 427 446,  
 449 50 437 438  
 Hook Device Service 43 44  
 Hook V86 Int Chan, 37  
 Hooks, 95 99 238 41  
 domain, 238 41  
 front end, 497 98  
 SHARE, 490 94

## I

i cur, 672  
 i select, 68 69  
 IBMBIO (C) 189 90 265, 266 454  
 IBMSYS 190  
 ICH, 396  
 ICMGR 496  
 IN WNTF flag 64  
 IN W NDRWS flag 6  
 IN 1110 : 242  
 INDRS End 35  
 INDRS : 358 39 365  
 InDOS flag 38 48 121 358 60  
 INI 566 IntApp 47  
 IntSys at 77  
 IntVADP1 162  
 Intell V86 Break Point 210  
 Intense data 24 30  
 INTR MID, 212 635 38 662  
 INI 2Ab network critical section 591  
 function 30b, 34 35 37 38  
 function 31b, 35 37 584 85  
 function 32b, 37 296 584 85  
 INT 21h multiplex 2 47 99 300  
 301 306 8  
 function 01b, 49 609 611  
 function 05b, 199 307 537  
 function 06b, 49  
 function 01b, 603  
 function 10b, 49 192 490 491  
 function 11b, 49 57 507 326  
 372 463 467 495 498 505 6  
 509 513 16, 519 23 526 534  
 function 12b, 10 68 89 190  
 279 294 300 301 306 14 384  
 500 518, 527 534 39 620  
 function 12b, 20 21 309  
 function 14b, 49  
 function 16b, 6 25 27 29 32 34  
 35 40 49 113 120 143 143,  
 157 162 173 238 564  
 function 13b, 49  
 function 27b, 186 192  
 function 22b, 548 608  
 function 22b, 237 248 268  
 function 40b, 41 210  
 function 43b, 278 70  
 function 4Ab, 40, 41 42 451 455  
 function 55b, 644  
 function 62b, 44  
 function 60b, 49  
 function 81b, 49  
 function 82b, 49  
 INT 21h, MS DPCs, See also Sys ar  
 INT 21b, function 57b  
 function 03b, 129  
 function 0Aa, 253 546, 564 623,  
 643 647  
 function 0Db, 198, 302 324 25  
 532, 564

function 0E, 60 61 68 69 93  
 95 96 98 98 99  
 function 0b, 231  
 function 0c, 199  
 function 0d, 25  
 function 0E, 198  
 function 10b, 68 69 264  
 function 10b, 176 7  
 function 10b, 126 27  
 function 12b, 198  
 function 15b, 48 49 126 27 324 26  
 function 24b, 198 499  
 function 25b, 125 128 199 279,  
 281 342 579  
 function 29b, 626 function 2Ab,  
 126 198  
 function 21b, 27  
 function 30b, 576 98 174  
 184 198 208  
 function 31b, 223 385 541 542 635  
 function 37b, 48 49 126 27 198  
 324 26 416 444 436  
 function 33b, 38 38 208 288 87  
 447  
 function 36b, 36 48 49, 26, 127  
 192 198 260 251 294, 324 28  
 541 548  
 function 35b, 198 271 279 542  
 function 37b, 198  
 function 38b, 14  
 function 39b, 174 74 74 24  
 function 3Ab, 74 4  
 function 3Bb, 522  
 function 3Cb, 238  
 function 31b, 20 124 24 217  
 74 74 48 58 58  
 function 3Db, 379 53  
 function 3Eb, 601 595  
 function 40b, 10, 127  
 function 47b, 326 29 431 338  
 389 316  
 function 43b, 185 186, 423  
 function 44b, 185 186, 48 198,  
 192 41, 420 436  
 function 46b, 626  
 function 47b, 174  
 function 48, 88 93 644 346  
 365 367  
 function 48b, 125 149  
 function 49b, 198 258 343  
 function 4Ab, 258 343 555 367  
 function 41b, 48 49 69 283  
 288 263 269 364 497  
 function 43b, 6 34, 605 613  
 function 41b, 185  
 function 42b, 128  
 function 50b, 68 48 135 97  
 199 750, 251 360 541 560  
 function 51b, 68 48 52 98  
 54, 300 360  
 function 52b, 198 384  
 function 54b, 7









- SEVER, 6, 74, 75 97, 184, 199
- SFIS (System File Tables), 2, 20, 21  
51, 92, 93, 180, 222
- CGUs and, 24
- IOFS emulation and, 208, 9
- IOFS file systems and, 430-438  
442, 465, 69
- IOFSIRG and, 27, 33
- IOFDOS and, 192
- KERNEL and, 37
- NetWare and, 196, 197
- MV5ASB1 (I) and, 44-45
- MFIF (F) and, 483, 85
- virtual machine and, 329, 30
- Windows and, 143, 44
- SETBACK, 130, 143, 48, 477  
file, 128, 41  
GPs to file and, 141, 13  
Windows and, 143, 44
- SETBACK, 21, 110, 11, 138, 469, 72
- SHARE, 38, 45, 192, 195, 202, 467  
books, 192, 490, 94
- The Phantom and, 508
- SHARE (S) 7K, 188, 192, 490, 94, 537
- SHARE (S) 3K, 491, 91
- SHARE (S) 3K, 492, 491
- SHIF - 373, 640, 664
- SHIF (2), 664, 65
- SHIF, 20
- Shell command, 224, 638, 41, 674  
show display, 104
- Silverberg, Brad, 16, 184
- SmartDrive, 39, 40, 441, 451
- SMARTDRY, 337, 8, 39
- SMARTDRY, 30
- Soft ICE/Windows, WFNDR, 9, 20  
111, 116, 117, 119  
prelink at IOFS bases and, 155  
Windows IOFS extenders and, 121
- Source, 41, 42
- spawn, 663
- SpeedLink, 43
- spintp, 659, 672
- SRE (S) 11R, 500
- srch att pr, 519
- SRE (S) 11K, 500
- SRE (S) 62, 76, 137
- STAB, 6, 554
- STAB, K ASM, 553, 556, 58, 565
- Stacker, 39, 42, 402, 449  
IOFS file systems and, 458, 61
- STAB (S) 11R, 41
- STDLIB, 76
- STDLIB, 620, 21, 629
- STDLIB, 466
- STDP system, 246
- strncpy, 396
- STR (S) 11R, 242
- SUBS, 38, 45, 214, 444, 446, 447  
465, 630
- suspend background, 6, 0
- suspend foreground, 6, 0
- Switching, 66d, 593, 96, 610, 627
- SWITCH (S) 11R, 67
- SWITCH (S) 11R, 593, 96
- Stokes, Alan,
- SYSTEM, 13, 68
- Sys Critical Int, 123
- SYSTEM, 36, 37
- System, 441, 475
- system, 663
- SYSTEM, 30, 55, 36, 37, 173  
IOFSIRG and, 124  
prelink at IOFS bases and, 160  
system, management and, 569, 70  
S12, 403, 111, 144, 477
- SHARE, 494
- SysVars (S) 21b, function, 52b, 2, 4  
179, 201, 250, 266, 350, 474, 624
- SARE, data structure and, 10  
accessing, 75, 76, IOFS, and, 74
- IOFS (S) and, 384  
disassembly and, 294, 299, 500, 124  
25, 327
- IOFS file systems and, 437, 441, 42
- IOFSIRG and, 21, 24, 30, 35, 126  
178
- DoubleSpace and, 42
- IOFS (S) and, 119, 129
- IOFS (S) and, 190, 191, 92
- nucleus and, 60
- GPs tasks and, 110, 112
- IOFS (S) and, 251, 53
- KERNEL and, 37
- IOFS (S) and, 69, 71, 2  
80, 84, 85, 87, 88, 93, 191
- MV5ASB1 (I) and, 44
- NETX and, 198
- OS, 2 IOFS settings and, 217  
prelink and, 106  
programming problems and, 45
- SETBACK and, 138
- SYSTEMS (S) and, 21, 22
- SYSTEMS (S) and, 106, 7  
utilities, 6, 38, 39
- WIN (S) and, 20, 21
- Windows and, 101, 106, 108, 112  
119, 21, 126, 29, 141
- Task Managers, 592, 600
- TEMP, 44
- TEMP, 607
- TEMP (S) 607, 8, 602
- TEMP (S) 603
- TempName, 395
- TEMPNAME, 395, 96
- TempName, 103
- TempName, 610, 11
- TEMP, 496
- TempName, 172
- TEMP (S) Turbo Pascal for Windows, 66, 83
- TEMP (S) ASM, 553
- TempName, 148, 51, 428, 30, 468
- TRUNCATE, 510, 514  
command interpreters and, 609, 638  
OS, 2 IOFS settings and, 212
- TRUNCATE, 149, 50
- TRUNCATE, 638, 41, 674
- TRUNCATE, 674
- TRUNCATE, 362
- TRUNCATE, 630
- TRUNCATE, 565, 567, 568, 78, 479,  
603, 609, 610
- TRUNCATE, 548, 555, 566, 604, 198  
FILE (S) and, 602, 604
- TRUNCATE, 78, 79
- TRUNCATE, 548, 605, 6
- TRUNCATE, Terminal and Sys. Resident files,  
tests, 37, 88, 99, 446, 448, 541, 47  
application wrappers and, 95  
command line arguments and, 503, 84  
errors, 7, 546
- TRUNCATE (S) 558, 65
- TRUNCATE and, 192  
extended linker, a reference, 640  
and, 361, 64  
generators, 548, 50, 565, 84, 603  
file interpreters and, 611  
implementation, 555, 96  
interrupt changing and, 552, 53, 601, 2  
- output, 280 and, 496, 65  
juggling stacks and, 556, 58  
keeping C programs resident and,  
553, 55
- keyboard interrupts and, 611  
to Microsoft and Borland C++,  
550, 56
- MS-DOS flags and, 538, 61
- multitasking and, 609, 18  
not going resident, 4, 555, 96  
printing and, 611  
renaming, 601, 3  
sample programs and, 614, 8  
Task Managers and, 602, 600  
timer interrupts and, 610, 11  
TRUNCATE and, 151  
Windows NT and, 223  
working with the IOFS (S), 584, 92
- TRUNCATE (S) 565, 567, 579, 83,  
603, 610
- Turbo Debugger, 116
- Turbo Pascal, 60, 61, 72, 82, 84  
IOFS calls from, 65, 67  
IOFS (S) and, 60  
"magic numbers" and, 72  
unconnected IOFS calls from, 82, 84  
Using arrangements, 3, 18, 45
- UNLINK, 38, 151, 475  
chain finding, 462, 64  
decoding IOFS subprograms and,  
558, 59  
determining string length and, 361



# About the Andrew Schulman Programming Series

---

**T**HE ANDREW SCHULMAN PROGRAMMING SERIES has been designed for PC programmers seeking clarity and direction in a rapidly evolving programming universe. This award-winning series reinforces the idea that a decade of programming experience in MS-DOS need not be abandoned in the brave new age of objects and windows.

All titles in the series are authoritative, timely, and substantive. They encompass every aspect of IBM programming with particular emphasis on Windows and DOS programming.

Andrew Schulman, the series editor, is a software engineer known for his insight on programming issues and his ability to express technical ideas clearly. He has established a reputation as one of the finest writers on programming topics.

## **Windows Internals**

*From the moment I opened this book, I was transformed from a DOS programmer to a Windows programmer. I have read this book several times and will continue to read it and keep it on the shelf for future reference.*

*Bob Mauer*

## **Undocumented DOS**

*This is a great book for anyone who is interested in the inner workings of DOS. It is a must-read for anyone who is a DOS programmer. I have read this book several times and will continue to read it and keep it on the shelf for future reference.*

*PC Master*

## **Undocumented Windows**

*Undocumented Windows is a must-read for anyone who is a Windows programmer. It is a great book for anyone who is interested in the inner workings of Windows. I have read this book several times and will continue to read it and keep it on the shelf for future reference.*

*PC Master*

## **Windows 3.1 Programming for Mere Mortals**

*This is a brilliant title and the book Windows 3.1 Programming for Mere Mortals is just a bass.*

*PC Techniques*

## **Windows ++**

*No other Windows book I have seen comes close to this. It is a must-read for anyone who is a C++ Windows programmer.*

*Dr. Hobbes*

## About the Books in the Series



### **Undocumented DOS: A Programmer's Guide to Reserved MS-DOS Functions and Data Structures, Second Edition**

By **ALAN R. FISHER** and **D. J. MILES**, *Raytheon*  
Microsoft's reserved DOS functions and data structures have become an integral addition to any serious programmer's library. This book provides a complete, up-to-date reference on these functions and structures, covering the latest DOS releases, including DOS 6.0, Windows™ 3.1, and DR-DOS™ 6.0, with coverage of all the new features introduced in the past year.

Microsoft's reserved DOS functions, it has become an integral addition to any serious programmer's library. This book provides a complete, up-to-date reference on these functions and structures, covering the latest DOS releases, including DOS 6.0, Windows™ 3.1, and DR-DOS™ 6.0, with coverage of all the new features introduced in the past year.

**\$44.95, Paperback, 880 pages, 3.5" disk**  
ISBN 0-201-63287-X



### **Windows ++: Writing Reusable Windows Code in C++**

By **Paul Dilavio** describes how to build a C++ class library. Rather than teach you how to use commercially available class libraries, it shows you how to build your own system, one that's tailored to suit your needs. Along the way, you'll learn how to use the Windows 3.11 API to create and manage windows, menus, and other graphical user interface elements. The book also shows you a number of various tips and techniques to help you get the most out of Windows 3.11.

**\$29.95, Paperback, 592 pages**  
ISBN 0-201-60891-X



### **Windows 3.1 Programming for Mere Mortals**

By **Woody Ziemhard** uses commonly available Windows 3.11 programming techniques to show you how to create a complete Windows 3.11 application from the ground up. You'll discover the power of creating dialog boxes, moving data between applications, and using the Windows 3.11 API to create and manage windows, menus, and other graphical user interface elements. The book also shows you a number of various tips and techniques to help you get the most out of Windows 3.11.

**\$44.95, Paperback, 560 pages, 3.5" disk**  
ISBN 0-201-60832-4



## Windows Internals: The Implementation of the Windows Operating Environment

Microsoft's Windows operating system is the most widely used operating system in the world. Windows Internals: The Implementation of the Windows Operating Environment is the definitive guide to the inner workings of Windows. This book is the most comprehensive and authoritative source of information on the Windows operating system. It covers the Windows kernel, the Windows user-mode architecture, and the Windows file system. It also covers the Windows networking stack, the Windows security architecture, and the Windows registry. This book is a must-read for anyone who is interested in the Windows operating system.

\$29.95, paperback 560 pages  
ISBN 0 201 62217 3



## The Undocumented PC: A Programmer's Guide to I/O, CPUs, and Fixed Memory Areas

by Frank van Gallewe reveals the undocumented or difficult-to-find information that is essential for the development of low-level software. This book covers the hardware and software details of the PC architecture, including the CPU, memory, and I/O. It also covers the undocumented features of the operating system, such as the BIOS, the boot process, and the hardware initialization. This book is a must-read for anyone who is interested in the low-level details of the PC architecture.

\$44.95 Paperback 864 pages 3.5" disk  
ISBN 0 201 62277 7



## DOS Internals

by Geoff K. Huppell is the definitive guide to exploiting the full potential of DOS. Meticulously researched, this book covers the inner workings of the operating system, including the kernel, the file system, and the hardware. It also covers the undocumented features of the operating system, such as the boot process, the hardware initialization, and the hardware initialization. This book is a must-read for anyone who is interested in the inner workings of the DOS operating system.

\$39.95 Paperback 688 pages 3.5" disk  
ISBN 0 201 60835 9




## **Undocumented Windows: A Programmer's Guide to Reserved Microsoft Windows API Functions**

by Andrew Schulman, David Macey, and Matt Pietrek is a complete guide and comprehensive reference to the Windows API functions left undocumented or "reserved" by Microsoft. The first section of the book introduces the inner workings of Windows and the role of the reserved API functions. The second section contains a comprehensive reference to all undocumented Windows functions. It names and defines each function, lists the versions and modes of Windows that support it, names programs and libraries currently using it, and then notes unique features, potential problems, and conflicts.

\$39.95, Paperback, 736 pages, 3.5" disk.

ISBN 0-201-60834-0



## **DOS and Windows Protected Mode: Programming with DOS Extenders in C**

by Al Williams is essential reading if you're using or contemplating using protected mode or DOS extender technology in applications development. It is the definitive guide, complete with practical source code, to writing applications that take full advantage of the most popular DOS extenders including Phar Lap and Intel. The accompanying disk contains a working model of Phar Lap's 286/DOS-Extender Lite.

\$39.95, Paperback, 528 pages, 5.25" disk.

ISBN 0-201-63218-7



## **Windows Network Programming: How to Survive in a World of Windows, DOS, and Networks**

by Ralph Davis is the first book to clearly address the key issues regarding Windows and networks. The book looks closely at the leading network standards, including: NetWare, Windows NT, Banyan Vines, and TCP/IP. It also develops a network-independent interface for Windows applications by determining what functionality should be standardized.

\$29.95, Paperback, 560 pages.

ISBN 0-201-58133-7

### **ORDER INFORMATION**

Available wherever computer books are sold or call Addison-Wesley at 1-800-358-4566 in the United States. Outside of the U.S. call your local Addison-Wesley office.

Addison-Wesley books are available at special discounts for bulk purchases by corporations, institutions, and other organizations. For more information, please contact our Corporate, Government and Special Sales Department at (617) 944-3700 x2431



# HUNGRY . . .

FOR MORE INSIDE INFORMATION ON WINDOWS AND DOS!

DO YOU WANT TO KNOW EVEN MORE . . .

about Windows and DOS undocumented features and functions, internal details, and other underground programming secrets . . . FROM TOP INDUSTRY AUTHORS!

DO YOU WANT TO QUALIFY TO WIN FREE BOOKS?

Get the latest information on Schulman Series books and occasional tips, techniques and insights from series authors. Tell us how to reach you and we'll send you, as available, series updates and words of wisdom from top industry programmers and writers. Send your name, address, any e-mail addresses, and comments or book ideas to:

Schulman Series  
Trade Computer Books Department  
Addison-Wesley Publishing Company  
1 Jacob Way  
Reading, MA 01867

CompuServe #74640,2405 (c/o Philip Sutherland)

PLUS . . . three times a year we'll pull a name from those we receive and send the lucky winner FREE two new Schulman Series titles!

## Attention 5 1/4" Disk Drive Users:

The disk to accompany *Undocumented DOS, Second Edition* is also available on two 5 1/4" high density disks. Please return the coupon below with a check for \$10.00 payable to Addison-Wesley to:

Addison-Wesley Publishing Company  
Order Department  
One Jacob Way  
Reading, MA 01867-9994

---

Please send me the two 5 1/4" disks (ISBN 0-201-40687-X) to accompany *Undocumented DOS, Second Edition*. I am enclosing a check for \$10.00.

Name \_\_\_\_\_  
Address \_\_\_\_\_  
City \_\_\_\_\_  
State \_\_\_\_\_  
Zip code \_\_\_\_\_

# UNDOCUMENTED DOS

  
 The Andrew Schulman  
 Programming Series



What the critics had to say about the first edition of *Undocumented DOS*:

"Only one or two books per year stand out as truly worthwhile efforts that we can use every day. *Undocumented DOS* is such a book. Serious DOS programmers should own a copy."

—Ray Duncan, *PC Magazine*

"*Undocumented DOS* is the most informative DOS programming book I have ever read."

—Al Stevens, *Dr. Dobbs' Journal*

*Undocumented DOS* is back and it's bigger and better than ever. The book that launched a thousand press releases returns to take on MS-DOS 5 and 6, Windows™ 3.1, and Novell DOS in the no-holds-barred "undocumented" tradition. This standard industry reference has been thoroughly rewritten and updated to include:

- Coverage of the DOS/Windows interface. Everything you (and the FTC!) wanted to know about how DOS and Windows work together: the DOS/MSGR interface, the Windows DOS extenders, DPMI, INT 2Fh calls provided by Windows, and much more
- MS-DOS 5 and 6, including Upper Memory Blocks, the Task Switcher, instance data, and DoubleSpace disk compression. Even includes a few peeks at DOS 7.0, its "Chicago"
- Coverage of other variants of DOS, such as Novell DOS, the OS/2 DOS box, DOS emulation under Windows NT, and NetWare changes to DOS
- All new coverage of disassembling DOS and DOS internals
- Expanded coverage of the Microsoft network redirector.

The powerful utilities disk includes:

- A new version of the popular INTSPY debugger
- PHANTOM, a network redirector SMS RAM disk
- DEVLOD for loading device drivers from the command line
- INTRLST, a complete reference of all DOS interrupts in an easily viewable electronic format.

**ANDREW SCHULMAN** is a contributing editor to *Microsoft Systems Journal* and *Dr. Dobbs' Journal* and is author of the Windows Source discursively linked from 3 Communications.

**DAVID NAXES** is cofounder of Smart Storage, Inc., a software company specializing in storage management systems. He previously managed the CUI Networks development project at Lotus.

**KELE BROWN**, an independent software developer, holds a Ph.D. in computer science from Carnegie Mellon University and is well known in the on-line community for maintaining the "Interrupt List."

**RODMOND J. MOORE** is a senior software engineer for Better Dickerson in Sparks, Maryland, and an independent consultant specializing in MS-DOS applications, system programs, and Microsoft Windows.

**JIM KYLL** has written extensively on MS-DOS and contributed to *The DOS Programmer's Reference* and *The MS-DOS Encyclopedia*.



9 780201 632873

ISBN 0-201-63287-X

**\$44.95 US**  
**\$57.95 CANADA**